CS 170 HW 14

Due on 2018-12-02, at 9:59 pm

1 (★) Study Group

List the names and SIDs of the members in your study group.

2 (—) Nostalgia

What's been your favorite homework problem this semester? Tell us the HW number and problem name/number, and briefly explain (a sentence or two) why you liked it.

3 (★★) Entanglement

$|\psi\rangle = \frac{1}{\sqrt{2}}|00\rangle + \frac{1}{\sqrt{2}}|11\rangle$ is one of the famous "Bell states," a highly entangled state of its two qubits. In this question we examine some of its strange properties.

(a) Suppose this Bell state could be decomposed as the (tensor) product of two qubits, the first state in $\alpha_0|0\rangle + \alpha_1|1\rangle$ and the second state in $\beta_0|0\rangle + \beta_1|1\rangle$. Write four equations that the amplitudes α_0, α_1, β_0, and β_1 must satisfy. Conclude that the Bell state cannot be decomposed.

(b) What is the result of measuring the first qubit of $|\psi\rangle$?

(c) What is the result of measuring the second qubit after measuring the first qubit?

(d) If the two qubits in state $|\psi\rangle$ are (physically) very far from each other, can you see why the answer to (c) is surprising?

Solution:

(a) We can conclude that $\alpha_0 * \beta_0 = \frac{1}{\sqrt{2}}$, $\alpha_0 * \beta_1 = \alpha_1 * \beta_0 = 0$, and $\alpha_1 * \beta_1 = \frac{1}{\sqrt{2}}$ which implies that one of $\alpha_0, \beta_1 = 0$. However, both $\alpha_0 * \beta_0$ and $\alpha_1 * \beta_1$ are not 0, which is a contradiction.

(b) $|0\rangle$ with probablity 1/2 and $|1\rangle$ with probability 1/2.

(c) After measuring the first qubit, the second qubit will have the exact same value.

(d) Even when separated by a large amount of distance, the two particles' states are fully dependent on each other.

4 (★★★★) Quantum Gates

(a) The Hadamard Gate acts on a single qubit and is represented by the following matrix:

$$\frac{1}{\sqrt{2}}\begin{bmatrix}1 & 1\\ 1 & -1\end{bmatrix}$$

Verify that this gate maps the basis states $|0\rangle$ and $|1\rangle$ to a superposition state that will yield 0 and 1 with equal probability, when measured. In other words, explicitly represent the bases as vectors, apply the gate as a matrix multiplication, and explain why the resulting vector will yield 0 and 1 with probabilities 1/2 each, when measured.

(b) Give a matrix representing a *NOT* gate. As in the previous part, explicitly show that applying your gate to the basis state $|0\rangle$ will yield the state $|1\rangle$ (and vice-versa).

(c) Give a matrix representing a gate that swaps two qubits. Explicitly show that applying this matrix to the basis state $|01\rangle$ will yield the state $|10\rangle$. Verify that this matrix is its own inverse.

Solution:

(a) The state $|1\rangle$ can be interpreted as the vector $[0,1]^T$. The results of applying the Hadamard gate to $|1\rangle$ is

$$\frac{1}{\sqrt{2}}\begin{bmatrix}1 & 1\\ 1 & -1\end{bmatrix}\begin{bmatrix}0\\ 1\end{bmatrix} = \begin{bmatrix}\frac{1}{\sqrt{2}}\\ -\frac{1}{\sqrt{2}}\end{bmatrix}$$

When measured, this yields 1 with probability $(\frac{1}{\sqrt{2}})^2 = \frac{1}{2}$ and yields 0 with probability $(\frac{-1}{\sqrt{2}})^2 = \frac{1}{2}$.

Likewise, $|0\rangle$ can be interpreted as the vector $[1,0]^T$, so applying the Hadamard gate gives

$$\frac{1}{\sqrt{2}}\begin{bmatrix}1 & 1\\ 1 & -1\end{bmatrix}\begin{bmatrix}1\\ 0\end{bmatrix} = \begin{bmatrix}\frac{1}{\sqrt{2}}\\ \frac{1}{\sqrt{2}}\end{bmatrix}$$

When measured, this yields 1 and 0 with probabilities $(\frac{1}{\sqrt{2}})^2 = \frac{1}{2}$, each.

(b) The following matrix represents a NOT gate:

$$\begin{bmatrix}0 & 1\\ 1 & 0\end{bmatrix}$$

To see this, we apply it to $|1\rangle$ to get:

$$\begin{bmatrix}0 & 1\\ 1 & 0\end{bmatrix}\begin{bmatrix}0\\ 1\end{bmatrix} = \begin{bmatrix}1\\ 0\end{bmatrix}$$

Likewise, applying it to $|0\rangle$ yields

$$\begin{bmatrix}0 & 1\\ 1 & 0\end{bmatrix}\begin{bmatrix}1\\ 0\end{bmatrix} = \begin{bmatrix}0\\ 1\end{bmatrix}$$

Thus, the gate is, in fact, a NOT gate.

(c) As in the book, We represent a two qubit state by $|\alpha\rangle = \alpha_{00}|00\rangle + \alpha_{01}|01\rangle + \alpha_{10}|10\rangle + \alpha_{11}|11\rangle$. This is the same as the vector $[\alpha_{00}, \alpha_{01}, \alpha_{10}, \alpha_{11}]^T$. Swapping the two qubits is equivalent to swapping the middle two values, since swapping two of the same values is unnecessary. This can be done by the following matrix:

$$\begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}$$

Applying this matrix to the basis state $|01\rangle$ yields

$$\begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} 0 \\ 1 \\ 0 \\ 0 \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ 1 \\ 0 \end{bmatrix}$$

Simple matrix multiplication will show that the matrix is its own inverse.

5 (★★) Multiplicative Weights

Recall from the notes that in the experts problem, if there are n experts and the best expert has cost m, the randomized multiplicative weights algorithm has expected cost at most $(1+\epsilon)m + \frac{\ln n}{\epsilon}$.

(a) We run the randomized multiplicative weights algorithm with two experts and believe the best expert will have cost 10000. What choice of ϵ should we use to minimize the bound on the cost of the algorithm?

(b) We run the randomized multiplicative weights algorithm with two experts. In all of the first 140 days, Expert 1 has cost 0 and Expert 2 has cost 1. If we chose $\epsilon = 0.01$, on the 141st day with what probability will we play Expert 1? (Hint: You can assume that $0.99^{70} = \frac{1}{2}$)

Solution:

(a) Taking the derivative of $(1+\epsilon)m + \frac{\ln n}{\epsilon}$ with respect to ϵ gives $m - \frac{\ln n}{\epsilon^2}$. Setting this equal to 0 and solving for ϵ gives $\epsilon = \sqrt{\frac{\ln n}{m}}$. Plugging in $n = 2, m = 10000$ gives $\epsilon = \sqrt{\frac{\ln 2}{10000}} \approx .008326$.

(b) The weight assigned to expert 1 is $.99^{0 \cdot 140} = 1$, while the weight assigned to expert 2 is $.99^{1 \cdot 140} \approx 1/4$. So, the probability we play expert 1 is $\frac{1}{1+1/4} = 4/5$.

6 (★★★★) Experts Alternatives

Recall the experts problem. Every day you must take the advice of one of n experts. At the end of each day t, if you take advice from expert i, the advice costs you some c_i^t in $[0,1]$. You want to minimize the regret R, defined as:

$$R = \frac{1}{T}(\sum_{t=1}^{T} c_{i(t)}^t - \min_{1 \le i \le n} \sum_{t=1}^{T} c_i^t)$$

where $i(t)$ is the expert you choose on day t. Your strategy will be probabilities where p_i^t denotes the probability with which you choose expert i on day t. Assume an all powerful adversary can look at your strategy ahead of time and decide the costs associated with each expert on each day. Give the maximum possible (expected) regret that the adversary can guarantee if your strategy is:

(a) Choose expert 1 at every step. That is, if $\forall t\ p_1^t = 1$ and C_i^t is the set of costs for all experts and all days, what is $\max_{C_i^t} R$?

(b) Any deterministic strategy. Note that a "deterministic strategy" can be thought of as a probability distribution that satisfies the following: $\forall t\ \exists i\ p_i^t = 1$.

(c) Always choose an expert according to some fixed probability distribution at every time step. That is, if for some $p_1 \ldots p_n$, $\forall t, p_i^t = p_i$, what is $\max_{C_i^t}(\mathbb{E}[R])$?

 What distribution minimizes the regret of this strategy? In other words, what is $\text{argmin}_{p_1 \ldots p_n} \max_{C_i^t}(\mathbb{E}[R])$?

This analysis should conclude that a good strategy for the problem must necessarily be randomized and adaptive.

Solution:

(a) 1. Consider the case where the cost of expert 1 is always 1 and the cost of expert 2 is always 0. Thus $C = \sum_{t=1}^{T}\sum_{i=1}^{n} p_i^t c_i^t = \sum_{t=1}^{T} c_1^t = T$, we have $C^* = \sum_{i=1}^{T} c_2^t = 0$, so regret is $\frac{1}{T}(T-0) = 1$.

(b) $\frac{n-1}{n}$. Consider the case where the cost of the chosen expert is always 1, and the cost of each other expert is 0. Let k be the least-frequently chosen expert, and let m_k be the number of times that expert is chosen. This will result in a regret of $\frac{1}{T}(T - m_k)$

Since the best expert is the one that is chosen least often, the best strategy will try to maximize the number of times we choose the expert that is chosen least often. This means we want to choose all the experts equally many times, so expert k is chosen in at most T/n of the rounds. Therefore, $m_k \le \frac{T}{n}$, thus the regret is at least $\frac{1}{T}(T - \frac{T}{n}) = \frac{n-1}{n}$

(c) $(1 - \min_i p_i)$. Like in part 1, the distribution is fixed across all days, so we know ahead of time which expert will be chosen least often in expectation. Let $k = \arg\min_i p_i$ be the expert with least cost. Let $c_k^t = 0$ for all t, and let $c_i^t = 1$ for all $i \neq k$ and for all t. This way, $C = T(1 - p_k)$, as this is a binomial random variable distributed as $C \sim Bin(T, 1-p_k)$. $C^* = 0$, so we end up with a regret of $\frac{1}{T}(C-C^*) = \frac{1}{T}(T(1-p_k)-0) = 1 - p_k$.

To minimize the expectation of R is the same as maximizing $\min_i p_i$. This is maximized by the uniform distribution, obtaining regret $\frac{n-1}{n}$ (this is the same worst case regret as in part 2).