Model counting for logical theories
Problem set 3

1. (Von Neumann's problem.) Suppose you are given a coin that turns up heads with probability $p \in (0, 1)$, but you don't know the value of p. Can you *simulate* a fair coin flip using just this (potentially biased) coin? (*Hint:* The running time of your algorithm may follow geometric distribution with parameter that depends on p.)

2. Suppose you are given a fair coin, i.e., one that turns up heads with probability $1/2$. Also suppose you are given a number $p \in (0, 1)$. Can you *simulate* a flip of a biased coin that turns up heads with probability p, just using the fair one?

3. Consider the following random experiment. There is a box with three balls, each of which can be either black or white. In a single step of the experiment, a ball is taken uniformly at random from the box and replaced with a (new) ball of the other colour. This step can be repeated arbitrarily many times.

   (a) Write the transition probability matrix for the Markov chain whose state is the number of black balls in the box.

   (b) Compute transition probabilities for two consecutive steps of the experiment.

   (c) Does the Markov chain have a stationary distribution? (*Hint:* Yes.)

   (d) Is the stationary distribution unique?

   (e) Does this Markov chain always converge to a stationary distribution? Why?

4. In the lecture on Thursday we claimed that the constant $\frac{3}{4}$ probability of success of the oracle $\mathcal{E}$ does not prevent us from using the oracle to implement an algorithm with confidence level $1 - \alpha$ for any error probability α. This exercise essentially asks you to demonstrate how this is possible.

   You are given an algorithm $\mathcal{A}$ for a decision problem L that is such that

$$x \in L \Longrightarrow \mathsf{P}[\,\mathcal{A}(x) \text{ accepts}\,] \geq 3/4$$
$$x \notin L \Longrightarrow \mathsf{P}[\,\mathcal{A}(x) \text{ accepts}\,] \leq 1/4.$$

   You want to design an algorithm $\mathcal{B}$ which

   - receives as input a vector $(x_1, \ldots, x_n)$ of inputs to $\mathcal{A}$, and $\alpha \in [0, 1]$,
   - returns a vector $(y_1, \ldots, y_n) \in \{0, 1\}^n$ such that with probability at least $1 - \alpha$ it holds for all $i \in \{1, \ldots, n\}$ that $y_i = 1$ iff $x_i \in L$.

   Algorithm $\mathcal{B}$ has a tunable parameter r and

   - runs r independent copies $A_{i,1}, \ldots, A_{i,r}$ of $\mathcal{A}$ on each x_i, and
   - returns a vector $(y_1, \ldots, y_n) \in \{0, 1\}^n$ where y_i is the result of the majority vote on x_i:

$$y_i = \begin{cases} 1 & \text{if } \sum_{j=1}^{r} \mathcal{A}_{i,j}(x_i)) \geq \frac{r}{2}, \\ 0 & \text{otheriwse.} \end{cases}$$

   For a given $\alpha \in [0, 1]$ and n, determine a value of the parameter r, for which with probability at least $1 - \alpha$ the the output of $\mathcal{B}$ meets the requirements.

5. Suppose you are given an implementation of the oracle $\mathcal{E}$ from the lecture on Thursday that guarantees for some fixed constants $0 < c < C$ that

$$\mathsf{mc}(\varphi) \geq C \cdot 2^m \Longrightarrow \mathsf{P}[\mathcal{E}(\varphi, n) = \mathsf{YES}] \geq \frac{3}{4}$$
$$\mathsf{mc}(\varphi) \leq c \cdot 2^m \Longrightarrow \mathsf{P}[\mathcal{E}(\varphi, n) = \mathsf{NO}] \geq \frac{3}{4}$$

How would you implement an oracle $\mathcal{E}'$ such that

$$\mathsf{mc}(\varphi) \geq 2^{m+1} \Longrightarrow \mathsf{P}\left[\mathcal{E}'(\varphi, n) = \mathsf{YES}\right] \geq \frac{3}{4}$$
$$\mathsf{mc}(\varphi) \leq 2^m \Longrightarrow \mathsf{P}\left[\mathcal{E}'(\varphi, n) = \mathsf{NO}\right] \geq \frac{3}{4}.$$

6. Show that random affine operators over the field of two elements form a family of pairwise independent hash functions. To do this, assume that $\boldsymbol{x}$ and $\boldsymbol{y}$ are distinct vectors from $\{0,1\}^n$, and $\boldsymbol{w}_1, \boldsymbol{w}_2$ are some vectors from $\{0,1\}^m$. Prove that if a matrix $A \in \{0,1\}^{m \times n}$ and a vector $\boldsymbol{b} \in \{0,1\}^m$ are chosen uniformly at random and independently, then the function $h\colon \{0,1\}^n \to \{0,1\}^m$ given by

$$h(\boldsymbol{x}) = A \cdot \boldsymbol{x} + \boldsymbol{b} \bmod 2$$

satisfies the equality $\mathsf{P}[\,h(\boldsymbol{x}) = \boldsymbol{w}_1, h(\boldsymbol{y}) = \boldsymbol{w}_2\,] = (1/2^m)^2$.
Hint: If X is any random variable and b is chosen uniformly at random from $\{0,1\}$ independently from X, then the sum $(X + b) \bmod 2$ has uniform distribution over $\{0,1\}$.

7. Complete the proof of the Leftover Hash Lemma (simplified version), which we have seen in the lecture on Thursday:
Let $\mathcal{H}$ be a family of pairwise independent hash functions $h\colon \{0,1\}^n \to \{0,1\}^m$. Let $S \subseteq \{0,1\}^n$ satisfy $|S| \geq 4/\rho^2 \cdot 2^m$ for some $\rho > 0$. For $h \in \mathcal{H}$, let Z be the cardinality of the set $\{w \in S\colon h(w) = 0^m\}$. Prove that

$$\mathsf{P}\left[\left|Z - \frac{|S|}{2^m}\right| \geq \rho \cdot \frac{|S|}{2^m}\right] \leq \frac{1}{4}.$$

8. Assuming the Leftover Hash Lemma (simplified version), show that the Estimate oracle that we have described has the required properties:

   (a) Prove that if $\mathsf{mc}(\varphi) \geq 1000 \cdot 2^m$, then the formula $\psi(\boldsymbol{x}) := \varphi(\boldsymbol{x}) \wedge (h(\boldsymbol{x}) = 0^m)$ is unsatisfiable with probability at most $1/4$.
   Guidelines: Suppose $Z = \mathsf{mc}(\psi)$. If $Z = 0$, then $|Z - 1000| \geq 1000$. Apply LHL and find a suitable parameter $\rho > 0$.

   (b) Prove that if $\mathsf{mc}(\varphi) \leq 0.001 \cdot 2^m$, then the formula $\psi(\boldsymbol{x}) := \varphi(\boldsymbol{x}) \wedge (h(\boldsymbol{x}) = 0^m)$ is satisfiable with probability at most $1/4$.
   Guidelines: Let S' be any set that contains $[\![\varphi]\!]$ and satisfies $|S'| \geq 4 \cdot 2^m/\rho^2$ for some $\rho > 0$, to be fixed later. Suppose $Z = \mathsf{mc}(\psi)$ and $Z' = \mathsf{mc}(\psi')$ where ψ' is some fixed formula that satisfies $[\![\psi']\!] = S'$. Show that if $Z \geq 1$, then $|Z' - 0.001| \geq 0.999$. Apply LHL to S' and find a suitable parameter $\rho > 0$.