SEEM5020 2026 Spring: Homework 2

This assignment is due at 1:30 pm, January 23, 2026. Please bring your solution and submit it at the beginning of the next class. Total weight of Homework 2: 4%.

- **Q1. [10 Point] Chernoff Bound (Lower Tail).** Let $X_1, X_2, \ldots, X_n$ be independent Bernoulli random variables where $\mathbb{P}[X_i = 1] = p_i$. Let $X = \sum_{i=1}^{n} X_i$ and $\mu = \mathbb{E}[X]$. Prove the lower tail of the Chernoff bound (Theorem 12 in Lecture 1 slides). Specifically, show that for any $0 < \delta < 1$:

$$\mathbb{P}[X \leq (1-\delta)\mu] \leq \left(\frac{e^{-\delta}}{(1-\delta)^{1-\delta}}\right)^{\mu} \leq e^{-\mu\delta^2/2}.$$

  Hint: Use the moment generating function technique similar to the upper tail proof. You may use the inequality $(1-\delta)\ln(1-\delta) + \delta \geq \delta^2/2$ *for* $\delta \in (0,1)$.

- **Q2. [10 Point] Hoeffding's Inequality (Additive Bound).** While the multiplicative Chernoff bound is powerful, sometimes additive error suffices. **Hoeffding's Inequality** utilizes the bounded range of random variables to provide a strong bound on the absolute deviation λ.

  Let $X_1, \ldots, X_n$ be independent Bernoulli random variables (can be generalized to continuous bounded r.v.). Let $X = \sum_{i=1}^{n} X_i$ and $\mu = \mathbb{E}[X]$. Prove the following additive bound (only prove the upper tail for simplicity):

$$\mathbb{P}[X \geq \mu + \lambda] \leq e^{-\frac{2\lambda^2}{n}}$$

  Hint: You may use Hoeffding's Lemma without proof: For a random variable Z *with* $\mathbb{E}[Z] = 0$ *and* $a \leq Z \leq b$, $\mathbb{E}[e^{tZ}] \leq e^{\frac{t^2(b-a)^2}{8}}$. *Apply this to the centered variables* $X_i - p_i$.

  Remark: A natural first thought is substituting $\delta = \lambda/\mu$ into the multiplicative Chernoff bound, and it will yield $e^{-\lambda^2/(3\mu)}$. In the TA's own experiences, Hoeffding's Inequality is applied more frequently than the Chernoff Bound.

- **Q3. [5 Point] Weighted Reservoir Sampling (A-Res).** In the A-Res Algorithm (Algorithm 2 in Lecture 2 slides), for each element (e_i, w_i), we generate a key $K_i = U_i^{1/w_i}$ where $U_i \sim \text{Uniform}(0,1)$, and maintain the k elements with the largest keys.

  Prove that this procedure is mathematically equivalent to generating keys $K_i' = -\frac{\ln(U_i)}{w_i}$ and maintaining the k elements with the **smallest** keys.

  Remark: The logarithmic form is often preferred in implementation for numerical stability to avoid underflow when w_i is small.

- **Q4. [15 Point] Morris(a) Algorithm.** Consider the Morris(a) algorithm (Exercise on Slide 21 of Lecture 2), where we increment the counter X with probability $\frac{1}{(1+a)^X}$. We estimate the count N as $\hat{N} = \frac{(1+a)^X - 1}{a}$.
  - I. [5 marks] Prove that $\hat{N}$ is an unbiased estimator of N, i.e., $\mathbb{E}[\hat{N}] = N$.
  - II. [5 marks] Prove that the variance of the estimator is $\text{Var}[\hat{N}] = \frac{aN(N-1)}{2}$.
  - III. [5 marks] Prove that with probability at least $1-\delta$, the number of bits required to store the counter X is $O(\log\log N + \log(1/\delta) + \log(1/\epsilon))$ by setting $a = 2\epsilon^2\delta$.