Exercise 1. Prove the other side of the Chernoff Bound. That is

$$Pr(X < (1-\delta)\mu) \leq \exp[-\delta^2\mu/2] \ .$$

Exercise 2. Consider a biased coin with the probability of getting heads being an unknown parameter p, which is known to be at least a, for some $a > 0$. A natural procedure for estimating the coin bias is to flip the coin n times, and estimate p as the fraction of times it lands on head. Denote this estimate by p' and suppose that for a given parameter we want to have

$$Pr[|p - p'| > \epsilon p] < \delta \ .$$

How many flips do we need in function of a, δ and ϵ ? Using a calculator/computer, compute this number for $a = 0.1, \epsilon = 0.1$ and $\delta = 0.01$.

Exercise 3. Suppose we are given a sequence of n distinct numbers $a_1, \ldots, a_n$ and we want to compute the median of the sequence. There is a deterministic linear time algorithm for this problem, but it uses $\Omega(n)$ memory which can be problematic for huge n.

The goal of this exercise is to analyze an algorithm that estimates the median using random sampling and whose required storage only depends on the quality of the estimate and not on n. The algorithm works as follows, select uniformly and independently at random k numbers of the sequence then output the median of those k numbers.

Define the rank of an element as its position in the sorted sequence. For instance the rank of the minimum, maximum and median element will be $1, n$ and $\lfloor n/2 \rfloor$ respectively.

Suppose that we want the sampling algorithm to return an element x whose rank is approximately $n/2$ with high probability. More precisely, we want

$$Pr\left[\frac{n}{2}(1-\epsilon) \leq \text{rank}(x) \leq \frac{n}{2}(1+\epsilon)\right] \geq 1 - \delta \ .$$

What k do we need to get the desired confidence?

Hint: Treat each inequality independently and try to express the probability in such a way that we can apply the Chernoff bound.