

Lecture 03: Concentration Bounds: Markov and Chebyshev Inequality

Definition: Expectation and Variance I

- Let \mathbb{X} be a random variable over the sample space $\Omega \subseteq \mathbb{R}$
- The “average value of \mathbb{X} ” (referred to as the *mean*) is defined as follows for a discrete sample space Ω

$$\mathbb{E}[\mathbb{X}] := \sum_{\omega \in \Omega} \omega \cdot \mathbb{P}[\mathbb{X} = \omega].$$

- Think: How to generalize this definition to continuous Ω , say $\Omega = \mathbb{R}$
- This definition of the mean coincides with the notion of the “center of gravity” of the probability mass function of the random variable \mathbb{X} .

Definition: Expectation and Variance II

- An essential property of the mean is the following

Theorem (Linearity of Expectation)

For two random variables \mathbb{X}, \mathbb{Y} over the sample space $\Omega \subseteq \mathbb{R}$, the following identity holds

$$\mathbb{E}[\mathbb{X} + \mathbb{Y}] = \mathbb{E}[\mathbb{X}] + \mathbb{E}[\mathbb{Y}].$$

This identity holds irrespective of whether \mathbb{X} and \mathbb{Y} are independent of each other or not!

- For example, let *Bernoulli distribution* $\text{Ber}(p)$ over the sample space $\{0, 1\}$ represent the distribution where $\mathbb{P}[\text{Ber}(p) = 1] = p$ and $\mathbb{P}[\text{Ber}(p) = 0] = (1 - p)$. Note that $\mathbb{E}[\text{Ber}(p)] = p$.

Definition: Expectation and Variance III

- Let $(\mathbb{X}_1, \dots, \mathbb{X}_n)$ be a joint distribution such that each marginal distribution \mathbb{X}_i is identical to the distribution $\text{Bern}(p)$. Then, by the linearity of expectation, we have $\mathbb{E} \left[\sum_{i=1}^n \mathbb{X}_i \right] = np$.
- The *variance* of a distribution (intuitively) is a measure of how tightly the distribution is distributed around its mean. It is defined as follows

$$\text{Var} [\mathbb{X}] := \mathbb{E} \left[(\mathbb{X} - \mu)^2 \right],$$

where $\mu = \mathbb{E} [\mathbb{X}]$.

- Prove: $\text{Var} [\mathbb{X}] = \mathbb{E} [\mathbb{X}^2] - \mu^2$.

Overview for Concentration Bounds

- Previously, we learned about the expected value of a random variable
- However, is the expected value a good representation of the random variable?
- If the random variable concentrates most of its probability mass around the expected value, then we consider the expected value to be a good representation of the random variable's behavior
- In the topics of concentration, we shall cover techniques to argue the “typicality of a randomized experiment,” i.e., say the mean (or the median) being a good representative of a probability distribution

Markov Inequality I

Theorem (Markov Inequality)

Let \mathbb{X} be a random variable over the sample space $\Omega \subseteq \mathbb{R}_{\geq 0}$ (i.e., the set of non-negative real numbers), and $\mu = \mathbb{E}[\mathbb{X}]$. Then, the following bound holds.

$$\mathbb{P}[\mathbb{X} \geq \lambda\mu] \leq \frac{1}{\lambda}$$

By substituting variables, this bound is also equivalent to the following expression.

$$\mathbb{P}[\mathbb{X} \geq \lambda] \leq \frac{\mu}{\lambda}$$

Intuition: Suppose λ is large. Then, the probability that \mathbb{X} deposits probability mass further than $\lambda\mu$ is unlikely.

I present the proof only for discrete Ω . The case of general Ω is similar.

Markov Inequality II

Proof. If possible, let the Markov inequality be false. That is, there exists $\lambda > 1$ such that $\mathbb{P}[X \geq \lambda\mu] > 1/\lambda$. Then, let us lower-bound the expectation as follows.

$$\begin{aligned}\mu = \mathbb{E}[X] &= \sum_{i \in \Omega} i \cdot \mathbb{P}[X = i] \\ &= \sum_{i \in \Omega: i < \lambda\mu} i \cdot \mathbb{P}[X = i] + \sum_{i \in \Omega: i \geq \lambda\mu} i \cdot \mathbb{P}[X = i] \\ &\geq \sum_{i \in \Omega: i < \lambda\mu} 0 \cdot \mathbb{P}[X = i] + \sum_{i \in \Omega: i \geq \lambda\mu} (\lambda\mu) \cdot \mathbb{P}[X = i] \\ &= 0 \cdot \mathbb{P}[X < \lambda\mu] + (\lambda\mu) \cdot \mathbb{P}[X \geq \lambda\mu] \\ &> (\lambda\mu) \cdot \frac{1}{\lambda} = \mu\end{aligned}$$

So, we have obtained a contradiction that $\mu > \mu$.

Comments on Markov I

- We emphasize that for every μ and $\lambda \geq 1$, there is a distribution for which the Markov inequality is tight. Let \mathbb{X} be a distribution such that $\mathbb{P}[\mathbb{X} = 0] = 1 - 1/\lambda$ and $\mathbb{P}[\mathbb{X} = \lambda\mu] = 1/\lambda$.
- If there exists B such that $\mathbb{P}[\mathbb{X} > B] = 0$, i.e., the sample space of \mathbb{X} is bounded above, then we can also apply Markov inequality to the random variable $(B - \mathbb{X})$
- Think: The pigeon-hole principle states that if m balls are placed arbitrarily into n bins, a bin with $\lceil m/n \rceil$ balls exists. How is Markov inequality equivalent to the pigeon-hole principle?

Comments on Markov II

- Think: Consider the following problem. Suppose (\mathbb{R}, \mathbb{C}) is a joint distribution over $\Omega = \{1, \dots, m\} \times \{1, \dots, n\}$. Intuitively, think of a matrix with m -rows and n -columns. The r.v. associates probability with the cells. Suppose there is a Fun event, and the following bound holds.

$$\mathbb{P}[(\mathbb{R}, \mathbb{C}) \in \text{Fun}] \geq \varepsilon$$

That is, if you sample a cell according to the joint distribution (\mathbb{R}, \mathbb{C}) , then the probability of the Fun event occurring is at least ε . Consider the following expression.

$$\mathbb{P}[(\mathbb{R}, \mathbb{C}) \in \text{Fun} | \mathbb{R} = r]$$

This expression represents the probability of the Fun event happening if we restrict (condition) on the row $r \in \{1, \dots, m\}$. Prove the following statement.

“The probability of sampling $r \sim \mathbb{R}$ such that it has

$$\mathbb{P}[(\mathbb{R}, \mathbb{C}) \in \text{Fun} | \mathbb{R} = r] \geq \alpha$$

is at least ε/α .”

Russel Impagliazzo refers to this result as the pigeon-hole principle. The proof of this result is similar to the proof of the Markov inequality. It is an excellent exercise to think of techniques to use this result for derandomization.

Chebyshev's Inequality I

Theorem (Chebyshev's Theorem)

For any real-valued random variable \mathbb{X} , the following bound holds

$$\mathbb{P} [|\mathbb{X} - \mu| \geq t] \leq \frac{\text{Var} [\mathbb{X}]}{t^2},$$

where $\mu = \mathbb{E} [\mathbb{X}]$.

Proof Outline.

$$\begin{aligned} \mathbb{P} [|\mathbb{X} - \mu| \geq t] &= \mathbb{P} [(\mathbb{X} - \mu)^2 \geq t^2] \\ &\leq \frac{\mathbb{E} [(\mathbb{X} - \mu)^2]}{t^2} \\ &= \frac{\text{Var} [\mathbb{X}]}{t^2} \end{aligned}$$

Chebyshev's Inequality II

- In the previous proof, we used the following fact.

$$\mathbb{P} [|\mathbb{X} - \mu| \geq t] = \mathbb{P} [(\mathbb{X} - \mu)^2 \geq t^2]$$

In general, this is true for any monotonically increasing function f . That is, for any monotonically increasing function $f: \mathbb{R} \rightarrow \mathbb{R}$, we have

$$\mathbb{P} [\mathbb{X} \geq t] = \mathbb{P} [f(\mathbb{X}) \geq f(t)]$$

This trick is extremely crucial and shall be used in various other problems.

- Additionally, the random variable $(\mathbb{X} - \mu)^2$ is non-negative. Therefore, we could apply the Markov inequality to conclude the following

$$\mathbb{P} [(\mathbb{X} - \mu)^2 \geq t^2] \leq \frac{\mathbb{E} [(\mathbb{X} - \mu)^2]}{t^2}$$

Chebyshev's Inequality III

- So, we saw that “Markov studied the r.v. X and got a bound in $1/t$ ” and “Chebyshev studied the r.v. X^2 and got a bound in $1/t^2$.” Can we extrapolate this approach to use “higher powers of X ” (technically referred to as the moments) to obtain bounds that are “higher polynomials in $1/t$?”
- Prove: For a random variable X over the sample space $\mathbb{R}_{\geq 0}$ the following identity holds.

$$\mathbb{P}[X \leq 0] \leq \frac{\text{Var}[X]}{(\mathbb{E}[X])^2}.$$

- Think: Find the random variable X over the sample space $[a, b]$ satisfying $\mathbb{E}[X] = \mu \in [a, b]$ that maximizes $\text{Var}[X]$.