

Lecture 22: Concentration Inequalities (Concluding Remarks)

Objective

- The objective of this lecture is to summarize all the concentration inequalities that we have studied in this course
- We shall also highlight some subtleties that differentiate the use of one concentration inequality from the others

Markov Inequality

- Let X be a random variable such that $X \geq 0$ (i.e., the random variable is non-negative)
- Then the only bound that we can claim is the following for any $t \geq 0$

$$\mathbb{P}[X \geq t] \leq \frac{\mathbb{E}[X]}{t}$$

- For example, this bound implies that if $t = \alpha \mathbb{E}[X]$, then $\mathbb{P}[X \geq t] \leq 1/\alpha$. That is, it is unlikely that the expectation of a random variable exceeds the expected value significantly
- **Comment:** For every t , there exists a random variable X such that the Markov inequality is tight. We emphasize that there is no single X that witnesses the tightness of the Markov inequality for all t

Chernoff Bound I

- Let X_i , for $1 \leq i \leq n$, be independent random variables. Each X_i is a random variable over the sample space $[0, 1]$.
- **Comment:** Note that the X_i s need not be identical. We can use “linear offset + linear scaling” to apply the Chernoff bound whenever there exists a, b such that $a \leq X_i \leq b$, for all $1 \leq i \leq n$
- Let $p = (\mathbb{E}[X_1] + \dots + \mathbb{E}[X_n])/n$
- Let $S_n := X_1 + \dots + X_n$
- Note that $\mathbb{E}[S_n] = np$. Our objective is to prove that the probability that S_n significantly exceeds its expected value is very small. Chernoff bound states the following.

$$\mathbb{P}[S_n \geq np + E] \leq \exp\left(-nD_{\text{KL}}(p + E/n, p)\right)$$

Chernoff Bound II

- **Comment:** The Chernoff bound is true for all p and n . For example, the value of p can depend on n itself. Consider $p = 1/\sqrt{n}$. The bound continues to hold
- We can apply the Chernoff bound to the sum of $1 - \mathbb{X}_i$ random variables to conclude that

$$\mathbb{P}[\mathbb{S}_n \leq np - E] \leq \exp\left(-nD_{\text{KL}}\left((1-p) + E/n, 1-p\right)\right)$$

- This form of the inequality, although very precise (tight lower-bound can be derived using the Stirling's Approximation and was part of your homework problems), is not easy to evaluate. This form is also unwieldy to understand the asymptotics based on n . So, we prove simpler to evaluate forms of this inequality

- Simpler Form 1.

$$\mathbb{P}[S_n \geq np + E] \leq \exp\left(-nD_{\text{KL}}\left(p + E/n, p\right)\right) \leq \exp(-2E^2/n)$$

This bound is very easy to evaluate and highlights that most of the probability mass is concentrated around $E \approx \sqrt{n}$ radius of concentration around the mean.

- **Comment:** Note that this bound is oblivious to the value of p . So, we can bound the lower tail as follows. Let $q = 1 - p$.

$$\mathbb{P}[S_n \leq np - E] \leq \exp\left(-nD_{\text{KL}}\left(q + E/n, q\right)\right) \leq \exp(-2E^2/n)$$

Chernoff Bound IV

- **Comment:** The property that the bound is independent of p is a significant drawback when we want to use Chernoff bound for problems where p is very small. For example, consider the case of $p = 1/\sqrt{n}$. In this case, the expected value $\mathbb{E}[S_n] = \sqrt{n}$, so a radius of concentration $E \approx \sqrt{n}$ is not very meaningful. So, we go for easier to evaluate multiplicative forms of Chernoff bounds.
- **Comment:** Easier to evaluate Form 1 of the Chernoff bound is a corollary of the Hoeffding's bound (which, in turn, is a corollary of the independent bounded difference inequality; which, in turn, is a corollary of the Azuma's inequality)

Chernoff Bound V

- Simpler Form 2. Let $\varepsilon > 0$ and $\mu = \mathbb{E}[S_n] = np$. Let $p' = p(1 + \varepsilon)$. We have

$$\mathbb{P}[S_n \geq \mu(1 + \varepsilon)] \leq \exp(-nD_{\text{KL}}(p', p)) \leq \exp\left(-\frac{\varepsilon^2 \mu}{2(1 + \frac{\varepsilon}{3})}\right)$$

This bound proves that the radius of concentration is roughly $E \approx \sqrt{\mu}$.

- Simpler Form 3. Let $1 > \varepsilon > 0$ and $\mu = \mathbb{E}[S_n] = np$. Let $q = 1 - p$ and $q' = 1 - p(1 - \varepsilon)$. We have

$$\mathbb{P}[S_n \leq \mu(1 - \varepsilon)] \leq \exp(-nD_{\text{KL}}(q', q)) \leq \exp\left(-\frac{\varepsilon^2 \mu}{2}\right)$$

This bound proves that the radius of concentration is roughly $E \approx \sqrt{\mu}$.

Hoeffding's Bound

- Let $a_i \leq X_i \leq b_i$ be independent random variables for $1 \leq i \leq n$. Define $c_i = b_i - a_i$, for $1 \leq i \leq n$
- Let $S_n := X_1 + \dots + X_n$
- Then, the following bound holds

$$\mathbb{P} [S_n \geq \mathbb{E} [S_n] + E] \leq \exp \left(-2E^2 / \sum_{i=1}^n c_i^2 \right)$$

- **Comment:** This bound is a consequence of the Bounded Independent Difference inequality
- **Comment:** The simpler to evaluate form 1 of Chernoff bound is a corollary of this bound

Independent Bounded Difference Inequality I

- Let X_1, \dots, X_n be independent random variables over the sample spaces $\Omega_1, \dots, \Omega_n$, respectively
- Define $\Omega := \Omega_1 \times \dots \times \Omega_n$
- Let $f: \Omega \rightarrow \mathbb{R}$ be any function that is (c_1, \dots, c_n) -bounded
- Let $\mu = \mathbb{E} [f(X_1, \dots, X_n)]$
- Then, the following bound holds

$$\mathbb{P} [f(X_1, \dots, X_n) \geq \mu + E] \leq \exp \left(-2E^2 / \sum_{i=1}^n c_i^2 \right)$$

- **Comment:** This bound is a consequence of the Azuma-Hoeffding inequality
- **Comment:** This bound yields Hoeffding's bound as a corollary

Independent Bounded Difference Inequality II

- **Comment:** The radius of concentration is $E \approx \sqrt{\sum_{i=1}^n c_i^2}$. Note that the “number of variables” also factors into this bound.
- **Comment:** We emphasize that the same “quantity” can be represented in multiple ways as functions of different random variables. For example, the chromatic number of a graph whose edges are uniformly and independently included in the graph with probability p . If we define a random variable for each edge, then there are $\binom{n}{2} \approx n^2/2$ random variables. The chromatic number is a $(1, \dots, 1)$ -bounded function.

Consequently, the radius of concentration is $\sqrt{\binom{n}{2}} \approx n$,

which is not meaningful (because $\mu = \Theta(n)$, for constant p).

Independent Bounded Difference Inequality III

However, if we consider the graph where \mathbb{X}_i represents how the vertex i is connected to the graph induced by $\{1, \dots, i-1\}$, then there are n random variables. The chromatic number is a $(0, 1, \dots, 1)$ -bounded function. So, we get a \sqrt{n} radius of concentration, which is meaningful.

- **Comment:** There are several applications of this bound. For example, to the chromatic number of random graphs, rank of random matrices
- **Comment:** The max-load function is also $(1, \dots, 1)$ -bounded. However, the radius of concentration is much-much larger than the expected value.

Azuma-Hoeffding Inequality I

- Let Ω be a sample space
- Consider the filtration

$$\{\emptyset, \Omega\} = \mathcal{F}_0 \subset \mathcal{F}_1 \subset \dots \subset \mathcal{F}_n$$

- Let $(\mathbb{F}_0, \mathbb{F}_1, \dots, \mathbb{F}_n)$ be a martingale sequence with respect to the filtration mentioned above
- For $1 \leq i \leq n$, there exists c_i such that for all $x \in \Omega$

$$c_i \geq \max_{y \in \mathcal{F}_{i-1}(x)} \mathbb{F}_i(y) - \min_{y \in \mathcal{F}_{i-1}(x)} \mathbb{F}_i(y)$$

- Then, the following bound holds

$$\mathbb{P}[\mathbb{F}_n \geq \mathbb{F}_0 + E] \leq \exp\left(-2E^2 / \sum_{i=1}^n c_i^2\right)$$

Azuma-Hoeffding Inequality II

- **Comment:** Let $\mathbb{X}_1 \in \Omega_1, \dots, \mathbb{X}_n \in \Omega_n$. Define $\Omega := \Omega_1 \times \dots \times \Omega_n$. Let $f: \Omega \rightarrow \mathbb{R}$ be any function. We emphasize that the distributions $\mathbb{X}_1, \dots, \mathbb{X}_n$ need not be independent. Suppose we are interested in studying the concentration of the function f around its expected value $\mu = \mathbb{E} [f(\mathbb{X}_1, \dots, \mathbb{X}_n)]$. That is, bound the quantity

$$\mathbb{P} [f(\mathbb{X}_1, \dots, \mathbb{X}_n) \geq \mu + E] \leq ?$$

- 1 We consider the filtration that reveals one coordinate of the functions at a time.
- 2 We consider the Doob's Martingale with respect to this filtration using the following random variables. For any $0 \leq i \leq n$ we define

$$\mathbb{F}_i(x) := \mathbb{E} [f(\mathbb{X}_1, \dots, \mathbb{X}_n) | \mathcal{F}_i] (x)$$

Note that $\mathbb{F}_0(x) = \mathbb{E} [f(\mathbb{X}_1, \dots, \mathbb{X}_n)] = \mu$.

Azuma-Hoeffding Inequality III

- 3 Now, we need to compute c_i , for $1 \leq i \leq n$, such that

$$c_i \geq \max_{y \in \mathcal{F}_{i-1}(x)} \mathbb{F}_i(y) - \min_{y \in \mathcal{F}_{i-1}(x)} \mathbb{F}_i(y)$$

We emphasize that $\max_{y \in \mathcal{F}_{i-1}(x)} \mathbb{F}_i(y)$ can be different from $\min_{y \in \mathcal{F}_{i-1}(x)} \mathbb{F}_i(y)$ for different $x \in \Omega$. However, as long as their difference is bounded by c_i for every $x \in \Omega$, we are all set!

- 4 Now, Azuma-Hoeffding inequality shall imply that

$$\mathbb{P}[\mathbb{F}_n \geq \mu + E] \leq \exp\left(-2E^2 / \sum_{i=1}^n c_i^2\right)$$

- **Comment:** We can obtain concentration bounds for (c_1, \dots, c_n) -bounded functions where each coordinate is uniformly and independently chosen.
- **Comment:** We can obtain the concentration of the Hypergeometric series and in the Pólya's urn experiment

Talagrand Inequality I

- Let $\mathbb{X} = (\mathbb{X}_1, \dots, \mathbb{X}_n)$ be independent random variables such that $\mathbb{X}_1 \in \Omega_1, \dots, \mathbb{X}_n \in \Omega_n$.
- Define the product space $\Omega := \Omega_1 \times \dots \times \Omega_n$
- Let $A \subseteq \Omega$ be any subset
- Talagrand's inequality states the following

$$\mathbb{P}[\mathbb{X} \in A] \cdot \mathbb{P}[d_{\mathcal{T}}(\mathbb{X}, A) \geq t] \leq \exp(-t^2/4)$$

Here $d_{\mathcal{T}}(\cdot, \cdot)$ represents the convex distance.

- **Comment:** Let us introduce the definition of an c -configuration function $f: \Omega \rightarrow \mathbb{R}$. The function f is c -configuration, if for every $x \in \Omega$, there exists α_x such that $\|\alpha_x\|_2 = 1$ and for every $y \in \Omega$ we have

$$f(y) \geq f(x) - \sqrt{cf(x)}d_{\alpha_x}(x, y)$$

Let m represent the median of the distribution $f(\mathbb{X})$. Then, the following bounds are corollary of the Talagrand's inequality.

$$\mathbb{P} [f(\mathbb{X}) \geq m + t] \leq \exp(-t^2/4c(m + t))$$

$$\mathbb{P} [f(\mathbb{X}) \leq m - t] \leq \exp(-t^2/4cm)$$

Talagrand Inequality III

- **Comment:** Remember that $m \leq 2\mathbb{E}[f(\mathbb{X})]$ by Markov inequality. This bound for c -configuration functions yields a concentration radius of $\sqrt{m} \approx \sqrt{\mu}$, which is applicable even when m is much less than n . However, note that we need $\mathbb{X}_1, \dots, \mathbb{X}_n$ to be independent of each other. So, Talagrand and Azuma-Hoeffding inequalities are incomparable to each other (that is, they do not subsume each other).
- **Comment:** In the class we saw an application to the longest increased subsequence of n random numbers in $[0, 1)$. The mean/median of this random variable is $\Theta(\sqrt{n})$. Talagrand's inequality yields a concentration of \sqrt{m} around the median.