

Chernoff Inequalities

Andreas Klappenecker

Texas A&M University

© 2018 by Andreas Klappenecker. All rights reserved.

Recall the following basic tail inequality.

Theorem (Markov's Inequality)

If X is a nonnegative random variable and t a positive real number, then

$$\Pr[X \geq t] \leq \frac{E[X]}{t}.$$

Proof.

Let Y denote the indicator random variable of the event $X \geq t$, so

$$Y(\omega) = \begin{cases} 1 & \text{if } X(\omega) \geq t, \\ 0 & \text{if } X(\omega) < t. \end{cases}$$

The expectation value of X satisfies

$$E[X] \geq E[tY] = t E[Y] = t \Pr[X \geq t],$$

which proves the claim. □

Definition

The **variance** $\text{Var}[X]$ of a discrete random variable X is defined by

$$\text{Var}[X] = E[(X - E[X])^2] = E[X^2] - E[X]^2,$$

whenever this expression is well-defined. The variance measures the squared deviation from the expected value $E[X]$.

As a consequence, we obtain the following concentration inequality.

Theorem (Chebychev's inequality)

If X is a random variable, then

$$\Pr[(X - E[X])^2 \geq \beta] \leq \frac{\text{Var}[X]}{\beta}. \quad (1)$$

Proof.

Given the random variable X , we can define the new random variables $Y = X - E[X]$ and Y^2 . Since Y^2 is a nonnegative random variable, Markov's inequality shows that

$$\Pr[Y^2 \geq \beta] \leq \frac{E[Y^2]}{\beta}.$$

Since $E[Y^2] = E[(X - E[X])^2] = \text{Var}(X)$, we have

$$\Pr[(X - E[X])^2 \geq \beta] = \Pr[Y^2 \geq \beta] \leq \frac{E[Y^2]}{\beta} = \frac{\text{Var}[X]}{\beta}. \quad \square$$

We can reformulate this in the following way.

Corollary (Chebychev's inequality)

If X is a random variable, then

$$\Pr[|X - E[X]| \geq t] \leq \frac{\text{Var}[X]}{t^2}. \quad (2)$$

Proof.

$$\Pr[|X - E[X]| \geq t] = \Pr[(X - E[X])^2 \geq t^2] \leq \frac{\text{Var}[X]}{t^2}. \quad \square$$

Consider now the sum of n independent random variables X_k ,

$$X = X_1 + X_2 + \cdots + X_n.$$

What kind of bound can we get in this case?

Can we improve on Chebychev's inequality in this case?

Sums of Independent Random Variables

Proposition

Let X_1, X_2, \dots, X_n be independent random variables. Then for real numbers $t > 0$, we get

$$\Pr \left[\left| \sum_{k=1}^n X_k - \sum_{k=1}^n E[X_k] \right| \geq t \right] \leq \frac{\sum_{k=1}^n \text{Var}[X_k]}{t^2}.$$

Proof.

We can apply Chebychev's inequality to $\sum_{k=1}^n X_k$. Since the random variables X_k are independent, the variance satisfies

$$\text{Var} \left[\sum_{k=1}^n X_k \right] = \sum_{k=1}^n \text{Var}[X_k]. \quad \square$$

Corollary

Let X_1, X_2, \dots, X_n be independent identically distributed random variables with $\mu = E[X_k]$. Then the average of these random variables satisfies

$$\Pr \left[\left| \frac{1}{n} \sum_{k=1}^n X_k - \mu \right| \geq t \right] \leq \frac{\text{Var}[X_1]}{nt^2}.$$

Corollary (Khinchin)

Let X_1, X_2, \dots be independent identically distributed random variables with $\mu = E[X_k]$. Then the average of these random variables satisfies

$$\lim_{n \rightarrow \infty} \Pr \left[\left| \frac{1}{n} \sum_{k=1}^n X_k - \mu \right| \geq t \right] = 0.$$

In other words, the sample average converges in probability towards the expected value μ .

Chernoff Bound

Consider now the sum of n independent bounded random variables X_k ,

$$X = X_1 + X_2 + \cdots + X_n,$$

What kind of bound can we get in this case?

Murphy's Law

Anything that can go wrong, does.

Let A_1, A_2, \dots, A_n denote mutually independent “bad” events. Murphy's law is the folk-wisdom that some bad event is likely to happen.

Let X_k denote the indicator random variable for the event A_k . So

$$\Pr[X_k = 1] = p_k = \Pr[A_k].$$

Let $X = \sum_{k=1}^n X_k$ denote number of bad events happening.

How likely is it that no “bad” event will happen?

Proposition

$$\Pr[X = 0] \leq e^{-E[X]}.$$

Proof.

$$\begin{aligned}\Pr[X = 0] &= \Pr[\overline{A_1 \cup A_2 \cup \dots \cup A_n}] \\ &= \Pr[\overline{A_1} \cap \overline{A_2} \cap \dots \cap \overline{A_n}] \\ &= \prod_{k=1}^n (1 - \Pr[A_k]) \\ &\leq \prod_{k=1}^n e^{-\Pr[A_k]} = e^{-\sum_{k=1}^n \Pr[A_k]} = e^{-\sum_{k=1}^n E[X_k]} = e^{-E[X]}. \quad \square\end{aligned}$$

Example

A production facility for micro processors boasts that the probability that a transistor will be faulty is merely 10^{-5} . You plan to produce a VLSI circuit with 10^6 transistors. Thus, the probability that you get a correctly working micro processor is less than

$$e^{-10} \approx 0.000045$$

The main reason why we obtained an improvement by going from Markov's inequality to Chebychev's inequality was that we considered a function of the random variable X .

We were able to use the higher moment X^2 to improve the accuracy in the bound.

If X is an arbitrary random variable, a and t real numbers with $t > 0$, then

$$\begin{aligned}\Pr[X \geq a] &= \Pr[e^{tX} \geq e^{ta}] \\ &\leq \frac{E[e^{tX}]}{e^{ta}}\end{aligned}$$

by Markov's inequality. We can now choose t to minimize the right-hand side.

The bounds depend on the random variables and on the value that we choose for t . There are many different Chernoff bounds as a result!

Proposition

Let X_1, X_2, \dots, X_n be independent random variables with values in the interval $[0, 1]$. If $X = X_1 + X_2 + \dots + X_n$ and $E[X] = \mu$, then for every $a > 0$ we get the bounds

- 1 $\Pr[X \geq \mu + a] \leq e^{-a^2/2n},$
- 2 $\Pr[X \leq \mu - a] \leq e^{-a^2/2n}.$

The random variables do not need to be Bernoulli random variables, but they need to be independent.

Proof.

Let $Y_k = X_k - E[X_k]$ for k in the range $1 \leq k \leq n$.

Then the mean $E[Y_k] = 0$.

Set $Y = Y_1 + Y_2 + \cdots + Y_n$. Then $Y = X - \mu$.

We can now apply the Bernstein's idea. For some $t > 0$, we get

$$\Pr[X \geq \mu + a] = \Pr[Y \geq a] = \Pr[e^{tY} \geq e^{ta}] \leq \frac{E[e^{tY}]}{e^{ta}}.$$

Proof. (Continued)

We have

$$\begin{aligned}\Pr[X \geq \mu + a] &\leq \frac{\mathbb{E}[e^{tY}]}{e^{ta}} = \frac{\mathbb{E}[e^{\sum_{k=1}^n tY_k}]}{e^{ta}} \\ &= \frac{\mathbb{E}[\prod_{k=1}^n e^{tY_k}]}{e^{ta}}.\end{aligned}$$

Since the random variables Y_k are independent, it follows that the random variables e^{tY_k} are independent. Therefore, we have

$$\Pr[X \geq \mu + a] \leq \frac{\mathbb{E}[\prod_{k=1}^n e^{tY_k}]}{e^{ta}} = e^{-at} \prod_{k=1}^n \mathbb{E}[e^{tY_k}]$$

Proof. (Continued)

It remains to bound $E[e^{tY_k}]$.

The function $f(y) = e^{ty}$ is convex, since $f''(y) = t^2 e^{ty} > 0$.

Let $c + dy$ be the line through the points $(-1, e^{-t})$ and $(1, e^t)$. So the coefficients c and d must satisfy

$$c = \frac{e^t + e^{-t}}{2} \quad \text{and} \quad d = \frac{e^t - e^{-t}}{2}.$$

By convexity of $f(y)$, we have

$$e^{ty} = f(y) \leq c + dy$$

for all y in $[-1, 1]$.

Proof. (Continued)

Therefore, we can conclude that

$$\begin{aligned} E[e^{tY_k}] &\leq E[c + dY_k] = c + d \underbrace{E[Y_k]}_{=0} \\ &= c = \frac{e^t + e^{-t}}{2}. \end{aligned}$$

Now we use the Taylor expansion of e^x to simplify the bound:

$$e^x = 1 + x + \frac{x^2}{2!} + \frac{x^3}{3!} + \dots$$

Proof. (Continued)

$$\begin{aligned}
E[e^{tY_k}] &\leq \frac{e^t + e^{-t}}{2} \\
&= \frac{1}{2} \left(1 + t + \frac{t^2}{2!} + \frac{t^3}{3!} + \dots \right) \\
&\quad + \frac{1}{2} \left(1 - t + \frac{t^2}{2!} - \frac{t^3}{3!} + \dots \right) \\
&= \left(1 + \frac{t^2}{2!} + \frac{t^4}{4!} + \frac{t^6}{6!} + \dots \right) \\
&\leq \left(1 + \frac{t^2}{2 \cdot 1!} + \frac{t^4}{2^2 \cdot 2!} + \frac{t^6}{2^3 \cdot 3!} + \dots \right) \quad \text{as } 2^k k! \leq (2k)! \\
&= e^{t^2/2}.
\end{aligned}$$

Proof. (Continued)

$$\Pr[X \geq \mu + a] \leq e^{-at} \prod_{k=1}^n \mathbb{E}[e^{tY_k}] \leq e^{-at} \prod_{k=1}^n e^{t^2/2} = e^{nt^2/2 - at}$$

The function $h(t) = n\frac{t^2}{2} - at$ has a minimum at $t_0 = a/n$. So

$$\Pr[X \geq \mu + a] \leq e^{n(a/n)^2 - a(a/n)} = e^{-a^2/(2n)}.$$

(Continued.)

We can obtain the second Chernoff inequality as follows.

Set $X' = -X$. Then $X \leq \mu - a$ if and only if $X' \geq -\mu + a$.

Therefore,

$$\Pr[X \leq \mu - a] \leq e^{-a^2/(2n)}.$$

This concludes the proof. □

Suppose that we toss a fair coin 10,000 times. Let $X_k = 1$ denote the event that the k -th coin toss yields heads. Let $X = X_1 + X_2 + \cdots + X_{10,000}$. Then

$$E[X] = 5000, \quad \text{Var}[X] = np(1 - p) = 2500.$$

① **Markov:** $\Pr[X \geq 6000] \leq 5000/6000 = 5/6$.

② **Chebychev:**

$$\Pr[X \geq 6000] = \Pr[X - E[X] \geq 1000] \leq 2500/10^6 = 1/400.$$

③ **Chernoff:**

$$\Pr[X \geq 6000] = \Pr[X \geq E[X] + 1000] \leq e^{-10^6/(2 \cdot 10,000)} = e^{-50}.$$

Proposition

Let $X = X_1 + X_2 + \cdots + X_n$ be the sum of n independent Bernoulli random variables with $\Pr[X_k = 1] = p_k$. Let

$$\mu = E[X] = p_1 + p_2 + \cdots + p_n.$$

Then

$$\Pr[X \geq (1 + \delta)\mu] \leq e^{\delta\mu}(1 + \delta)^{-(1+\delta)\mu},$$

$$\Pr[X \leq (1 - \delta)\mu] \leq e^{\delta\mu}(1 - \delta)^{-(1-\delta)\mu}.$$

These are good bounds, but they are rarely used in this form. We will find more convenient, but looser bounds, as a consequence.

Proof.

For all $t > 1$, we have

$$\begin{aligned}\Pr[X \geq (1 + \delta)\mu] &= \Pr[t^X \geq t^{(1+\delta)\mu}] \\ &\leq \frac{\mathbb{E}[t^X]}{t^{(1+\delta)\mu}} \quad \text{by Markov's inequality} \\ &= \frac{\mathbb{E}[t^{X_1} t^{X_2} \dots t^{X_n}]}{t^{(1+\delta)\mu}} \\ &= \frac{\mathbb{E}[t^{X_1}] \mathbb{E}[t^{X_2}] \dots \mathbb{E}[t^{X_n}]}{t^{(1+\delta)\mu}} \quad \text{by independence}\end{aligned}$$

For all $t > 1$, we have

$$\begin{aligned}\Pr[X \geq (1 + \delta)\mu] &= \frac{\prod_{k=1}^n \mathbb{E}[t^{X_k}]}{t^{(1+\delta)\mu}} \\ &= \frac{\prod_{k=1}^n ((1 - p_k) + p_k t)}{t^{(1+\delta)\mu}} \\ &= \frac{\prod_{k=1}^n (1 + p_k(t - 1))}{t^{(1+\delta)\mu}} \\ &\leq \frac{\prod_{k=1}^n e^{p_k(t-1)}}{t^{(1+\delta)\mu}} \\ &= \frac{e^{(t-1) \sum_{k=1}^n p_k}}{t^{(1+\delta)\mu}} = \frac{e^{(t-1)\mu}}{t^{(1+\delta)\mu}}\end{aligned}$$

Proof. (Continued)

We showed that for all $t \geq 1$, we have

$$\Pr[X \geq (1 + \delta)\mu] \leq \frac{e^{(t-1)\mu}}{t^{(1+\delta)\mu}}.$$

Substituting $t = 1 + \delta$ yields

$$\Pr[X \geq (1 + \delta)\mu] \leq \frac{e^{\delta\mu}}{(1 + \delta)^{(1+\delta)\mu}},$$

which is our claim.

The proof of the second inequality is done in a similar way. \square

Corollary

$$\Pr[X \geq (1 + \delta)\mu] \leq e^{-\delta^2\mu/3} \quad \text{for } 0 < \delta < 1,$$

$$\Pr[X \leq (1 - \delta)\mu] \leq e^{-\delta^2\mu/2} \quad \text{for } 0 < \delta < 1.$$

Proof.

For the first inequality, it suffices to show that

$$e^{\delta}(1 + \delta)^{-(1+\delta)} \leq e^{-\delta^2/3}$$

for $0 < \delta < 1$. Taking logarithms on both sides yields the equivalent inequality

$$f(\delta) := \delta - (1 + \delta) \ln(1 + \delta) + \frac{\delta^2}{3} \leq 0$$

The latter inequality is not hard to show. Indeed,

$$f'(\delta) = 1 - \frac{1 + \delta}{1 + \delta} - \ln(1 + \delta) + \frac{2}{3}\delta = -\ln(1 + \delta) + \frac{2}{3}\delta.$$

Proof.

$$f'(\delta) = -\ln(1 + \delta) + \frac{2}{3}\delta.$$

Since $f'(\delta) \leq 0$ for all δ in the range $0 \leq \delta \leq 1$, we can conclude that $f(\delta)$ is decreasing on $0 \leq \delta \leq 1$. Since $f(0) = 0$, it follows that

$$f(\delta) = \delta - (1 + \delta) \ln(1 + \delta) + \frac{\delta^2}{3} \leq 0.$$

Therefore, we can conclude that

$$\Pr[X \geq (1 + \delta)\mu] \leq e^{-\delta^2\mu/3} \quad \text{for } 0 < \delta < 1$$

holds. The proof of the second inequality is similar. \square

Probability Amplification

Recall that following definition.

Definition

Let ϵ be a constant in the range $0 \leq \epsilon < 1/2$.

The class **BPP** consists of all languages L such that there exists a polynomial-time randomized algorithm A such that

- 1 $x \in L$ implies $\Pr[A(x) \text{ accepts}] \geq 1 - \epsilon$,
- 2 $x \notin L$ implies $\Pr[A(x) \text{ rejects}] \geq 1 - \epsilon$.

It is customary to choose $\epsilon = 1/3$.

Let A be a randomized algorithm that decides $L \in \mathbf{BPP}$.

Let us construct an algorithm A' that runs A on an input x precisely n times and returns the majority vote as an answer.

How likely is it that A' errs?

Probability Amplification

Let X_k denote the indicator random variable that the k -th run returns the correct result. Let

$$X = \sum_{k=1}^n X_k$$

denote the number of runs with correct answer. Then

$$E[X] \geq \frac{2}{3}n.$$

Then A' gives the wrong answer with probability

$$\Pr \left[\sum_{k=1}^n X_k \leq \frac{n}{2} \right] \leq \Pr \left[\sum_{k=1}^n X_k \leq E[X] - \frac{n}{6} \right] \leq e^{-(n/6)^2/(2n)} = e^{-n/72}.$$

Therefore, if we choose $n = 72 \ln(1/\delta)$, then the algorithm errs with probability δ .

One can reduce the constant 72 by using other versions of the Chernoff bound.

However, the point is that the error δ of A' can be made as small as we please.

There are many different Chernoff bounds. The recipe to bound the tail of random variable X is as follows. If a and t real numbers with $t > 0$, then

$$\begin{aligned}\Pr[X \geq a] &= \Pr[e^{tX} \geq e^{ta}] \\ &\leq \frac{E[e^{tX}]}{e^{ta}}\end{aligned}$$

by Markov's inequality. We can now choose t to minimize the right-hand side.

The right-hand side is often bounded in various ways so that the bound is easier to use. All the resulting bounds are referred to as Chernoff bounds. You can find somewhat tighter bounds and variations in our textbook.

M. Mitzenmacher, E. Upfal, Probability and Computing, 2nd edition, 2017, Chapter 4.

Stasys Jukna, *Crashkurs Mathematik für Informatiker*, B.G. Teubner, 2008.