

6.262: Discrete Stochastic Processes 2/9/11

Lecture 3: Laws of large numbers, convergence

Outline:

- Review of probability models
- Markov, Chebychev, Chernoff bounds
- Weak law of large numbers and convergence
- Central limit theorem and convergence
- Convergence with probability 1

1

Review of probability models

Probability models are natural for real-world situations that are repeatable, using trials that

- have the same initial conditions
- are essentially isolated from each other
- have a fixed set of possible outcomes
- have essentially 'random' individual outcomes.

For any model, an extended model for a sequence or an n -tuple of IID repetitions is well-defined.

Relative frequencies and sample averages (in the extended model) 'become deterministic' and can be compared with real-world relative frequencies and sample averages in the repeated experiment.

2

The laws of large numbers (LLN's) specify what 'become deterministic' means.

They only operate within the extended model, but provide our only truly experimental way to compare the model with repeated trials of the real-world experiment.

Probability theory provides many many consistency checks and ways to avoid constant experimentation.

Common sense, knowledge of the real-world system, focus on critical issues, etc. often make repeated trials unnecessary.

The determinism in large numbers of trials underlies much of the value of probability.

3


Markov, Chebychev, Chernoff bounds

Inequalities, or bounds, play an unusually large role in probability. Part of the reason is their frequent use in limit theorems and part is an inherent imprecision in probability applications.

One of the simplest and most useful bounds is the Markov inequality: If Y is a non-negative rv with an expectation $E[Y]$, then for any real $y > 0$,


$$\Pr\{Y \geq y\} \leq \frac{E[Y]}{y}$$

Pf:



4

Markov inequality: $\Pr\{Y \geq y\} \leq \frac{E[Y]}{y}$



Note that the Markov bound is usually very loose. It is tight (satisfied with equality) if Y is binary with possible values 0 and y .

The Markov bound decreases very slowly (as $1/y$) with increasing y .

5

The Chebyshev inequality: If Z has a mean $E[Z] = \bar{Z}$ and a variance, σ_Z^2 , then for any $\epsilon > 0$,

$$\Pr\{|Z - \bar{Z}| \geq \epsilon\} \leq \frac{\sigma_Z^2}{\epsilon^2} \quad (1)$$

Pf: Let $Y = (Z - \bar{Z})^2$. Then $E[Y] = \sigma_Z^2$ and for any $y > 0$,

$$\Pr\{Y \geq y\} \leq \sigma_Z^2/y; \quad \Pr\{\sqrt{Y} \geq \sqrt{y}\} \leq \sigma_Z^2/y$$

Now $\sqrt{Y} = |Z - \bar{Z}|$. Setting $\epsilon = \sqrt{y}$ yields (1).

Chebyshev requires a variance, but decreases as $1/\epsilon^2$ with increasing distance ϵ from the mean.

6

The Chernoff bound: For any $z > 0$ and any $r > 0$ such that the moment generating function $g_Z(r) = E[e^{rZ}]$ exists,

$$\Pr\{Z \geq z\} \leq g_Z(r) \exp(-rz) \quad (2)$$

Pf: Let $Y = e^{rZ}$. Then $E[Y] = g_Z(r)$. For any $y > 0$, Markov says,

$$\Pr\{Y \geq y\} \leq g_Z(r)/y; \quad \Pr\{e^{rZ} \geq e^{rz}\} \leq g_Z(r)/e^{rz},$$

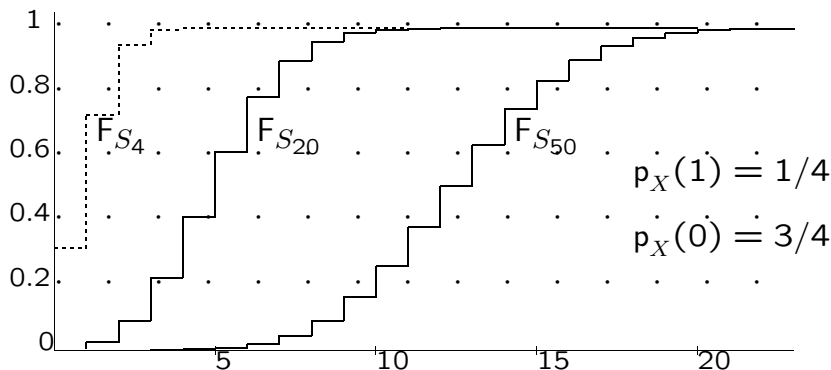
which is equivalent to (2).

This decreases exponentially with z and is useful in studying large deviations from the mean.

7

The weak law of large numbers and convergence


Let X_1, X_2, \dots, X_n be IID rv's with mean \bar{X} , variance σ^2 . Let $S_n = X_1 + \dots + X_n$. Then $\sigma_{S_n}^2 = n\sigma^2$.



The mean of the distribution varies with n and the standard deviation varies with \sqrt{n} .

8

The sample average is S_n/n , which is a rv of mean \bar{X} and variance σ^2/n .



The mean of the distribution is \bar{X} and the standard deviation decreases with $1/\sqrt{n}$.

9

$$\text{VAR} \left(\frac{S_n}{n} \right) = E \left[\left(\frac{S_n}{n} - \bar{X} \right)^2 \right] = \frac{\sigma^2}{n}. \quad (3)$$

$$\lim_{n \rightarrow \infty} E \left[\left(\frac{S_n}{n} - \bar{X} \right)^2 \right] = 0. \quad (4)$$

Note that (3) says more than (4), since it says the convergence is as $1/n$ and in fact it gives the variance explicitly. But (4) establishes a standard form of convergence of rv's called convergence in mean square.

Def: A sequence of rv's, Y_1, Y_2, \dots converges in mean square to a rv Y if

$$\lim_{n \rightarrow \infty} E \left[(Y_n - Y)^2 \right] = 0$$

10

The fact that S_n/n converges in mean square to \bar{X} doesn't tell us directly what might be more interesting: what is the probability that $|S_n/n - \bar{X}|$ exceeds ϵ as a function of ϵ and n ?

Applying Chebyshev to (3), however,

$$\Pr\left\{\left|\frac{S_n}{n} - \bar{X}\right| \geq \epsilon\right\} \leq \frac{\sigma^2}{n\epsilon^2} \quad \text{for every } \epsilon > 0 \quad (5)$$

One can get an arbitrary accuracy of ϵ between sample average and mean with probability $1 - \sigma^2/n\epsilon^2$, which can be made as close to 1 as we wish, by increasing n .

This gives us the weak law of large numbers (WLLN):

$$\lim_{n \rightarrow \infty} \Pr\left\{\left|\frac{S_n}{n} - \bar{X}\right| \geq \epsilon\right\} = 0 \quad \text{for every } \epsilon > 0.$$

11

WLLN: $\lim_{n \rightarrow \infty} \Pr\left\{\left|\frac{S_n}{n} - \bar{X}\right| \geq \epsilon\right\} = 0 \quad \text{for every } \epsilon > 0.$

We have proven this under the assumption that $S_n = \sum_{n=1}^n X_n$ where X_1, X_2, \dots , are IID with finite variance.

An equivalent statement (following from the definition of a limit of real numbers) is that for every $\delta > 0$,


$$\Pr\left\{\left|\frac{S_n}{n} - \bar{X}\right| \geq \epsilon\right\} \leq \delta \quad \text{for all large enough } n. \quad (6)$$

Note that (6) tells us less about the speed of convergence than

$$\Pr\left\{\left|\frac{S_n}{n} - \bar{X}\right| \geq \epsilon\right\} \leq \frac{\sigma^2}{n\epsilon^2}$$

But (6) holds without a variance (if $E[|X|] < \infty$.)

12



What this says is that $\Pr\left\{\frac{S_n}{n} \leq x\right\}$ is approaching a unit step at \bar{X} as $n \rightarrow \infty$. For any fixed ϵ , δ goes to 0 as $n \rightarrow \infty$. If $\sigma_X < \infty$, then $\delta \rightarrow 0$ at least as $\sigma^2/n\epsilon^2$. Otherwise it might go to 0 more slowly.

Def: A sequence of rv's, Y_1, Y_2, \dots converges in probability to a rv Y if for every $\epsilon > 0, \delta > 0$,

$$\Pr\{|Y_n - Y| \geq \epsilon\} \leq \delta \quad \text{for all large enough } n$$

This means that $\{S_n/n; n \geq 1\}$ converges to \bar{X} in probability if $E[|X|] < \infty$.

13

Review: We saw that if σ_X exists and X_1, X_2, \dots are IID, then $\sigma_{S_n/n} = \sigma_X/\sqrt{n}$.

Thus S_n/n converges to \bar{X} in mean square. Chebychev then shows that S_n/n converges to \bar{X} in probability.

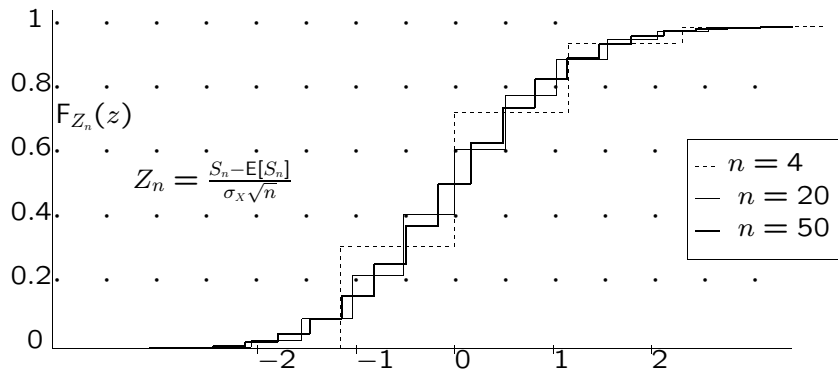
In the same way, if $\{Y_n; n \geq 1\}$ converges to Y in mean square, Chebychev show that it converges in probability.

That is, mean square convergence implies convergence in probability. The reverse is not true, since a variance is not required for the WLLN.

Finally, convergence in probability means that the distribution of $Y_n - Y$ approaches a unit step at 0.

14

Recall that $S_n - n\bar{X}$ is a zero mean rv with variance $n\sigma^2$. Thus $\frac{S_n - n\bar{X}}{\sqrt{n}\sigma}$ is zero mean, unit variance.



Central limit theorem:

$$\lim_{n \rightarrow \infty} \left[\Pr \left\{ \frac{S_n - n\bar{X}}{\sqrt{n}\sigma} \leq y \right\} \right] = \int_{-\infty}^y \frac{1}{\sqrt{2\pi}} \exp \left(\frac{-x^2}{2} \right) dx.$$

15

$$\lim_{n \rightarrow \infty} \left[\Pr \left\{ \frac{S_n - n\bar{X}}{\sqrt{n}\sigma} \leq y \right\} \right] = \int_{-\infty}^y \frac{1}{\sqrt{2\pi}} \exp \left(\frac{-x^2}{2} \right) dx.$$

Not only does $(S_n - n\bar{X})/\sqrt{n}\sigma_X$ have mean 0, variance 1 for all n , but it also becomes normal Gaussian.

We saw this for the Bernoulli case, but the general case is messy and the proof (by Fourier transforms) is not insightful.

The CLT applies to F_{S_n} , not to the PMF or PDF.

Def: A sequence Z_1, Z_2, \dots of rv's converges in distribution to Z if $\lim_{n \rightarrow \infty} F_{Z_n}(z) = F_Z(z)$ for all z where $F_Z(z)$ is continuous.

The CLT says that $(S_n - n\bar{X})/\sqrt{n}\sigma_X$ converges in distribution to Φ .

16

Convergence in distribution is almost a misnomer, since the rv's themselves do not necessarily become close to each other in any ordinary sense.

For example any sequence of IID rv's converge in distribution since they have the same distribution to start with.

Thm: Convergence in probability implies convergence in distribution.

Pf: Convergence of $\{Y_n; n \geq 1\}$ in probability means convergence to a unit step.

Thus convergence in mean square implies convergence in probability implies convergence in distribution.

17

Paradox: The CLT says something very strong about how S_n/n converges to \bar{X} , but convergence in distribution is a very weak form of convergence.

Resolution: The rv's that converge in distribution in the CLT are $(S_n - n\bar{X})/\sqrt{n}\sigma_X$. Those that converge in probability to 0 are $(S_n - n\bar{X})/n$, a squashed version of $(S_n - n\bar{X})/\sqrt{n}\sigma_X$.

The CLT, for $0 < \sigma_X < \infty$, for example, says that $\lim_{n \rightarrow \infty} \Pr\{(S_n - n\bar{X})/n \leq 0\} = 1/2$. This can not be deduced from the WLLN.

18

Convergence with probability 1

A rv is a far more complicated thing than a number. Thus it is not surprising that there are many types of convergence of a sequence of rv's.

A very important type is convergence with probability 1 (WP1). We introduce convergence WP1 here and discuss it more in Chap. 4.

The definition is deceptively simple.

Def: A sequence Z_1, Z_2, \dots , of rv's converges WP1 to a rv Z if

$$\Pr\{\omega \in \Omega : \lim_{n \rightarrow \infty} Z_n(\omega) = Z(\omega)\} = 1$$

19

$$\Pr\{\omega \in \Omega : \lim_{n \rightarrow \infty} Z_n(\omega) = Z(\omega)\} = 1$$

In order to parse this, note that each sample point maps into a sequence of real numbers, $Z_1(\omega), Z_2(\omega), \dots$.

Some of those sequences of real numbers have a limit, and in some cases, that limit is $Z(\omega)$. Convergence WP1 means that the set ω for which $Z_1(\omega), Z_2(\omega)$ has a limit, and that limit is $Z(\omega)$, is an event and that the probability of that event is 1.

One small piece of complexity that can be avoided here is looking at the sequence $\{Y_i = Z_i - Z; i \geq 1\}$ and asking if that sequence converges to 0 WP1.

20

The strong law of large numbers (SLLN) says the following: Let X_1, X_2, \dots be IID rv's with $E[|X|] < \infty$. Then $\{S_n/n; n \geq 1\}$ converges to \bar{X} WP1. In other words, all sample paths of $\{S_n/n; n \geq 1\}$ converge to \bar{X} except for a set of probability 0.

These are the same conditions under which the WLLN holds. We will see, when we study renewal processes, that the SLLN is considerably easier to work with than the WLLN.

It will take some investment of time to feel at home with the SLLN, (and in particular to have a real sense about these sets of probability 1) and we put that off until chap 4.

MIT OpenCourseWare
<http://ocw.mit.edu>

6.262 Discrete Stochastic Processes
Spring 2011

For information about citing these materials or our Terms of Use, visit: <http://ocw.mit.edu/terms>.