Weak Law of Large Numbers & Central Limit Theorem

Math 30530, Fall 2013

December 10, 2013

Weak Law of Large Numbers

$X_1, X_2, X_3, \ldots, X_n$ are independent copies of the same random variable, all with mean μ

Weak Law of Large Numbers

$X_1, X_2, X_3, \ldots, X_n$ are independent copies of the same random variable, all with mean μ

$M_n = \frac{X_1 + X_2 + \ldots + X_n}{n}$

Weak Law of Large Numbers

$X_1, X_2, X_3, \ldots, X_n$ are independent copies of the same random variable, all with mean μ

$M_n = \frac{X_1+X_2+...+X_n}{n}$

Weak law of large numbers: For every $\varepsilon > 0$,

$$\Pr(|M_n - \mu| \geq \varepsilon) \to 0 \text{ as } n \to \infty$$

Weak Law of Large Numbers

$X_1, X_2, X_3, \ldots, X_n$ are independent copies of the same random variable, all with mean μ

$M_n = \frac{X_1+X_2+\ldots+X_n}{n}$

Weak law of large numbers: For every $\varepsilon > 0$,

$$\Pr(|M_n - \mu| \geq \varepsilon) \to 0 \text{ as } n \to \infty$$

Interpretation: Repeat an experiment many times independently, record the answers, and average them. By making "many" large enough, we can be sure that the average of the readings is arbitrarily close to the theoretical experiment average, with arbitrarily high probability (arbitrarily close to 1)

Effective version

$X_1, X_2, X_3, \ldots, X_n$ are independent copies of the same random variable, all with mean μ, variance σ^2, $M_n = \frac{X_1+X_2+\ldots+X_n}{n}$

Effective version

$X_1, X_2, X_3, \ldots, X_n$ are independent copies of the same random variable, all with mean μ, variance σ^2, $M_n = \frac{X_1+X_2+\ldots+X_n}{n}$

Weak law of large numbers: For every $\varepsilon > 0$,

$$\Pr(|M_n - \mu| \geq \varepsilon) \leq \frac{\sigma^2}{n\varepsilon^2}$$

Effective version

$X_1, X_2, X_3, \ldots, X_n$ are independent copies of the same random variable, all with mean μ, variance σ^2, $M_n = \frac{X_1+X_2+\ldots+X_n}{n}$

Weak law of large numbers: For every $\varepsilon > 0$,

$$\Pr(|M_n - \mu| \geq \varepsilon) \leq \frac{\sigma^2}{n\varepsilon^2}$$

Example: I roll s dice 1,000 times. How sure can I be that the average of the rolls is between 3 and 4?

Effective version

$X_1, X_2, X_3, \ldots, X_n$ are independent copies of the same random variable, all with mean μ, variance σ^2, $M_n = \frac{X_1+X_2+\ldots+X_n}{n}$

Weak law of large numbers: For every $\varepsilon > 0$,

$$\Pr(|M_n - \mu| \geq \varepsilon) \leq \frac{\sigma^2}{n\varepsilon^2}$$

Example: I roll s dice 1,000 times. How sure can I be that the average of the rolls is between 3 and 4?

Solution: Here X_i is result of roll of dice, $\mu = 3.5$, $\sigma^2 = 2.92$, $n = 1000$, $\varepsilon = .5$, and

$$\Pr(|M_{1000} - 3.5| \geq .5) \leq \frac{2.92}{1000(.5)^2} = .01168,$$

so I can be around 98.8% sure of an average between 3 and 4.

More involved example

Example: A certain brand of lightbulb has lifetime that is exponentially distributed with mean A hours, A unknown. I try to estimate A by letting n lightbulbs run independently, and recording & averaging their lifetimes. How large should n be, so that I can be at least 90% sure that the estimate I get is within 5% of the actual average A?

More involved example

Example: A certain brand of lightbulb has lifetime that is exponentially distributed with mean A hours, A unknown. I try to estimate A by letting n lightbulbs run independently, and recording & averaging their lifetimes. How large should n be, so that I can be at least 90% sure that the estimate I get is within 5% of the actual average A?

We'll use $\Pr(|M_n - \mu| \geq \varepsilon) \leq \sigma^2/(n\varepsilon^2)$

More involved example

Example: A certain brand of lightbulb has lifetime that is exponentially distributed with mean A hours, A unknown. I try to estimate A by letting n lightbulbs run independently, and recording & averaging their lifetimes. How large should n be, so that I can be at least 90% sure that the estimate I get is within 5% of the actual average A?

We'll use $\Pr(|M_n - \mu| \geq \varepsilon) \leq \sigma^2/(n\varepsilon^2)$

Here $M_n = (X_1 + \ldots + X_n)/n$ with $X_i \sim \mathrm{exponential}(\lambda)$ (λ unknown), $\mu = 1/\lambda = A$, $\sigma^2 = 1/\lambda^2 = A^2$, $\varepsilon = .05\mu = .05A$, so

$$\Pr(|M_n - A| \geq .05A) \leq \frac{A^2}{n(.05A)^2} = \frac{400}{n}$$

Want n large enough so that this probability is at most .1, so $n = 4000$ large enough.

Another example — accumulating rounding error

Example: I estimate the sum of n random real numbers by rounding each to the nearest integer, and adding the resulting integers. What is the probability that the total error is at most $\pm\sqrt{n}$?

Another example — accumulating rounding error

Example: I estimate the sum of n random real numbers by rounding each to the nearest integer, and adding the resulting integers. What is the probability that the total error is at most $\pm\sqrt{n}$?

Let X_i be error in ith number, so $X_i \sim \mathrm{Uniform}(-1/2, 1/2)$. Let T_n be the total error, so $T_n = X_1 + \ldots + X_n$, all X_i independent. Each X_i has $\mu = 0$, $\sigma^2 = 1/12$. By weak law of large numbers,

$$\Pr(|M_n - 0| \geq \varepsilon) \leq \frac{1}{12n\varepsilon^2},$$

so, since $T_n = nM_n$, $\Pr(|T_n| \geq n\varepsilon) \leq \frac{1}{12n\varepsilon^2}$. Want $\Pr(|T_n| \geq \sqrt{n})$, so set $\varepsilon = 1/\sqrt{n}$, to get

$$\Pr(|T_n| \geq \sqrt{n}) \leq \frac{(\sqrt{n}^2)}{12n} = \frac{1}{12}.$$

Example: with 400 numbers, I can be at least $11/12$ sure that the error I make using rounding is no more than ± 20.

Central Limit Theorem

Informally: Add together lots of independent copies of a random variable. The result is very close to being a normal random variable.

Central Limit Theorem

Informally: Add together lots of independent copies of a random variable. The result is very close to being a normal random variable.

Less informally: Let $X_1, X_2, \ldots, X_n$ be a collection of independent copies of the same random variable, with X_i having mean μ and variance σ^2. Then

$$X_1 + \ldots + X_n \approx \text{Normal}(n\mu, n\sigma^2).$$

Central Limit Theorem

Informally: Add together lots of independent copies of a random variable. The result is very close to being a normal random variable.

Less informally: Let $X_1, X_2, \ldots, X_n$ be a collection of independent copies of the same random variable, with X_i having mean μ and variance σ^2. Then

$$X_1 + \ldots + X_n \approx \mathrm{Normal}(n\mu, n\sigma^2).$$

Precisely: Let $X_1, X_2, \ldots, X_n$ be as above. Set

$$S_n = \frac{X_1 + \ldots + X_n - n\mu}{\sqrt{n}\sigma}$$

(so $E(S_n) = 0$, $\mathrm{Var}(S_n) = 1$). Then, for each $x \in \mathbb{R}$,

$$\Pr(S_n \leq x) \to \Pr(Z \leq x) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{x} e^{-z^2/2}\, dz$$

as $n \to \infty$.

Example — accumulating rounding error

Example: I estimate the sum of n random real numbers by rounding each to the nearest integer, and adding the resulting integers. What is the probability that the total error is at most $\pm\sqrt{n}$?

Example — accumulating rounding error

Example: I estimate the sum of n random real numbers by rounding each to the nearest integer, and adding the resulting integers. What is the probability that the total error is at most $\pm\sqrt{n}$?

Let X_i be error in ith number, so $X_i \sim \mathrm{Uniform}(-1/2, 1/2)$. Let T_n be the total error, so $T_n = X_1 + \ldots + X_n$, all X_i independent. Each X_i has $\mu = 0$, $\sigma^2 = 1/12$. By central limit theorem, $T_n \approx \mathrm{Normal}(0, n/12)$, so

$$\begin{aligned}\Pr(|T_n| \le \sqrt{n}) &\approx \Pr(-\sqrt{n} \le \mathrm{Normal}(0, n/12) \le \sqrt{n}) \\ &= \Pr(-\sqrt{12} \le Z \le \sqrt{12}) \\ &\approx .99946.\end{aligned}$$

Example: with 400 numbers, I can actually be at least .999 sure that the error I make using rounding is no more than ± 20 (and at least .95 sure of no more than ± 11.3 error).