

Central Limit Theorem

The Theory

The central limit theorem proves that the **sum** of *any* iid random variables is normally distributed (as number of variables tends towards infinity). Consider IID random variables X_1, X_2, \dots such that $E[X_i] = \mu$ and $\text{Var}(X_i) = \sigma^2$ (since the variables are identically distributed μ and σ^2 is the same for all X_i). Let Y be the sum of the X_i s:

$$Y = \sum_{i=1}^n X_i$$

The central limit theorem states:

$$Y \sim N(n\mu, n\sigma^2) \quad \text{as } n \rightarrow \infty$$

Big deal! At this point you probably think that the central limit theorem is awesome. But it gets even better. With some algebraic manipulation we can show that if the **sum** of IID random variables is normal, it follows that the **mean** of equally weighted IID random variables must also be normal. Let's call the mean of IID random variables \bar{X} :

$$\bar{X} = \frac{1}{n} \sum_{i=1}^n X_i = \frac{1}{n} \cdot Y$$

If we define Y to be the sum of our variables

$$\sim N\left(\mu, \frac{\sigma^2}{n}\right)$$

Since Y is a normal and $\frac{1}{n}$ is a constant.

In summary, the central limit theorem explains that the sum of IID variables is normal (regardless of what distribution the IID variables came from) and the mean of equally weighted IID random variables is normal (again, regardless of the underlying distribution).

Example 1

Say you have a new algorithm and you want to test its running time. You have an idea of the variance of the algorithm's run time: $\sigma^2 = 4\text{sec}^2$ but you want to estimate the mean: $\mu = t\text{sec}$. You can run the algorithm repeatedly (IID trials). How many trials do you have to run so that your estimated runtime = $t \pm 0.5$ with 95% certainty? Let X_i be the run time of the i -th run (for $1 \leq i \leq n$).

$$0.95 = P\left(-0.5 \leq \frac{\sum_{i=1}^n X_i}{n} - t \leq 0.5\right)$$

By the central limit theorem, the standard normal Z must be equal to:

$$\begin{aligned} Z &= \frac{(\sum_{i=1}^n X_i) - n\mu}{\sigma\sqrt{n}} \\ &= \frac{(\sum_{i=1}^n X_i) - nt}{2\sqrt{n}} \end{aligned}$$

Now we rewrite our probability inequality so that the central term is Z :

$$\begin{aligned}
 0.95 &= P\left(-0.5 \leq \frac{\sum_{i=1}^n X_i}{n} - t \leq 0.5\right) = P\left(\frac{-0.5\sqrt{n}}{2} \leq \frac{\sum_{i=1}^n X_i}{n} - t \leq \frac{0.5\sqrt{n}}{2}\right) \\
 &= P\left(\frac{-0.5\sqrt{n}}{2} \leq \frac{\sqrt{n}}{2} \frac{\sum_{i=1}^n X_i}{n} - \frac{\sqrt{n}}{2} t \leq \frac{0.5\sqrt{n}}{2}\right) = P\left(\frac{-0.5\sqrt{n}}{2} \leq \frac{\sum_{i=1}^n X_i}{2\sqrt{n}} - \frac{\sqrt{n}}{2} \frac{\sqrt{nt}}{2} \leq \frac{0.5\sqrt{n}}{2}\right) \\
 &= P\left(\frac{-0.5\sqrt{n}}{2} \leq \frac{\sum_{i=1}^n X_i - nt}{2\sqrt{n}} \leq \frac{0.5\sqrt{n}}{2}\right) \\
 &= P\left(\frac{-0.5\sqrt{n}}{2} \leq Z \leq \frac{0.5\sqrt{n}}{2}\right)
 \end{aligned}$$

And now we can find the value of n that makes this equation hold.

$$\begin{aligned}
 0.95 &= \phi\left(\frac{\sqrt{n}}{4}\right) - \phi\left(-\frac{\sqrt{n}}{4}\right) = \phi\left(\frac{\sqrt{n}}{4}\right) - (1 - \phi\left(\frac{\sqrt{n}}{4}\right)) \\
 &= 2\phi\left(\frac{\sqrt{n}}{4}\right) - 1 \\
 0.975 &= \phi\left(\frac{\sqrt{n}}{4}\right) \\
 \phi^{-1}(0.975) &= \frac{\sqrt{n}}{4} \\
 1.96 &= \frac{\sqrt{n}}{4} \\
 n &= 61.4
 \end{aligned}$$

Thus it takes 62 runs. If you are interested in how this extends to cases where the variance is unknown, look into variations of the students' t-test.

Example 2

You will roll a 6 sided dice 10 times. Let X be the total value of all 10 dice $= X_1 + X_2 + \dots + X_{10}$. You win the game if $X \leq 25$ or $X \geq 45$. Use the central limit theorem to calculate the probability that you win.

Recall that $E[X_i] = 3.5$ and $\text{Var}(X_i) = \frac{35}{12}$.

$$\begin{aligned}
 P(X \leq 25 \text{ or } X \geq 45) &= 1 - P(25.5 \leq X \leq 44.5) \\
 &= 1 - P\left(\frac{25.5 - 10(3.5)}{\sqrt{35/12}\sqrt{10}} \leq \frac{X - 10(3.5)}{\sqrt{35/12}\sqrt{10}} \leq \frac{44.5 - 10(3.5)}{\sqrt{35/12}\sqrt{10}}\right) \\
 &\approx 1 - (2\phi(1.76) - 1) \approx 2(1 - 0.9608) = 0.0784
 \end{aligned}$$