Hypothesis testing and the Gamma function

Kenneth A. Ribet

Math 10A
November 29, 2016

Announcements

Please send me email to sign up for the next two (last two?) breakfasts:

- This Thursday (December 1) at 9AM.
- Next Monday (December 5), also at 9AM.

Pop-in lunch on Friday, December 2 at 12:10PM (no sign-ups needed, just show up)

Coffee at Blue Bottle Coffee Company (University Avenue near Li Ka Shing) on Friday, December 2 at 9AM.

Today we will discuss the Z-test (through examples), meet the Gamma function and (probably) start our discussion of Student's t-testing.

The Z-test business is sort of a review of last week's class.

The t-subject is the last topic in the Math 10A syllabus.

On Thursday, we will discuss t distributions and testing, maybe take a class photo and then have a course evaluation session where I leave and the students fill out their course evaluations online while they're still in the room.

This comes from the math department:

> *If your class meets on Tuesdays and Thursdays you could mention on Tuesday that you will be devoting the last 15 minutes of Thursday's class to the online course evaluations, and you could ask the students to bring their laptop computers to Thursdays class. . . . It would be a good idea for you to leave the classroom for those last 15 minutes while students are completing their evaluations of your teaching. Students can access the evaluation system through the following link:* `https://course-evaluations.berkeley.edu`.

As of 1PM on Tuesday, about 100 students had completed their Math 10A evaluations.

Last Thursday's pop-in lunch (non-vegetarian version)

Z-test example

Wikipedia (rephrased):

> *Suppose that in a particular geographic region, the mean and standard deviation of reading test scores are 100 points, and 12 points, respectively. . . . Fifty-five students in School X . . . received a mean score of 96. . . . are the students in this school comparable to a random sample of 55 students from the region as a whole, or are their scores surprisingly low?*

The scores of random generic students are given as values of a random variable X with mean $\mu = 100$ and standard deviation $\sigma = 12$. If we have N independent copies X_i of X (with $N = 55$, for example), then

$$\overline{X} = \frac{X_1 + \cdots + X_N}{N}$$

is (approximately) normally distributed with mean μ and standard deviation $\dfrac{\sigma}{\sqrt{N}}$.

As we all saw last Tuesday,

$$Z = \frac{(\overline{X} - \mu)\sqrt{N}}{\sigma}$$

is then normal(ish) with mean 0 and standard deviation 1.

An essentially correct statement is that Z is a standard normal variable in the limit (i.e., for $N \to \infty$).

When we evaluate this random variable on our sample of students, $\overline{X} = 96$ (the mean score of the 55 students in question) and $N = 55$. Thus

$$Z = \frac{(96 - 100)\sqrt{55}}{12} \approx -2.472.$$

This value occurs far out from the mean (which is 0). The probability of getting a Z that's -2.472 or less can be bounded by using the table on page 566 of Schreiber.

Namely, 49.18% of the scores occur between 0 and 2.4; thus fewer than 1% of the scores are larger than 2.4. Symmetrically, fewer than 1% of the scores are smaller than -2.4. Thus the probability of getting -2.472 or a lower score is way less than 1%. (Wikipedia says it's about 0.0068.)

If you prefer two-sided statements, you can say that fewer than 2% of students in a random sample will get adjusted scores that are either greater than 2.472 or less than -2.472. In any case, we're discussing *p*-values (probabilities) that are less than 5%.

Conclusion? We *reject* the null hypothesis that the 55 students from the particular school were drawn from a random sample of students in the entire region.

Other ways to communicate this information:

- The students from this school have surprisingly low scores.
- There is statistical significance to the low scores of the students.

Second example: Example 5.2 from Prob-Stat online notes

"We see a man flipping a coin and he gets 65 heads out of 100 tosses. If p is the probability of heads, we wish to test whether $p \neq 0.5$."

Why is the coin flipper a man?

This example is like that in the HW where 100 tosses led to 70 heads. In that situation, we rejected the null hypothesis that the coin is fair. This is again the null hypothesis, but 65 is less than 70, so the situation is less clear-cut.

If we plug in the data, we get for

$$\frac{(\overline{X} - \mu)\sqrt{N}}{\sigma}$$

the numerical value

$$\frac{(0.65 - 0.5) \cdot 10}{1/2} = 0.15 \cdot 20 = 3.$$

Being three standard deviations away from the mean is an extremely low-probability event. (It's worse than being 2.4 standard deviations away from the mean.) Specifically, being three standard deviations *above* the mean occurs with probability $1 - 0.9987$ (according to page 566 of the book). This number is about 0.13%. Being at least three standard deviations away from the mean (i.e., three above or three below) occurs with double the probability. The number involved, about .26%, is still minuscule.

We reject the null hypothesis. We think that the coin is not fair.

A Penn State example

Boys of a certain age have a mean weight of 85 pounds. A complaint is made that the boys living in a municipal children's home are underfed. Twenty-five boys (of this age) are weighed and found to have a mean weight of 80.94 pounds. The population standard deviation is 11.6 pounds. . . . what should be concluded concerning the complaint?

The null hypothesis is that the mean weight of children in this home is 85 pounds. The statistic Z is $\dfrac{80.94 - 85}{11.6}\sqrt{25} = -1.75$. For a normal distribution with mean 0 and standard deviation 1, the probability of being ≤ -1.75 is $0.5 - 0.4599$, according to table 7.3 on page 566 of Schreiber. This number is around 0.04, which is less than 0.05. Hence we reject the null hypothesis and conclude that the complaint is valid.

I WILL PASS MY FINALS.
I WILL PASS MY CLASSES.
I WILL NOT BE DEFEATED.
I WILL FINISH STRONG.

In quite a few situations involving statistics, one encounters values of the Gamma function $\Gamma(s)$. Here is a quick tutorial on Γ.

Definition: for real numbers $s > 0$,

$$\Gamma(s) = \int_0^\infty t^s e^{-t} \frac{dt}{t} = \int_0^\infty t^{s-1} e^{-t}\, dt.$$

For example,

$$\Gamma(1) = \int_0^\infty e^{-t}\, dt = -e^{-t} \Big]_0^\infty = 1.$$

Integration by parts:

$$\Gamma(s) = \int_0^\infty t^{s-1}e^{-t}\,dt = -t^{s-1}e^{-t}\,\Big]_0^\infty + \int_0^\infty (s-1)t^{s-2}e^{-t}\,dt$$
$$= 0 + (s-1)\Gamma(s-1).$$

Thus, for example:

$$\Gamma(2) = 1\cdot\Gamma(1), \quad \Gamma(3) = 2\cdot\Gamma(2), \quad \Gamma(4) = 3\cdot\Gamma(3), \ldots.$$

Since $\Gamma(1) = 1$, we get $\Gamma(2) = 1$, $\Gamma(3) = 2$, $\Gamma(4) = 6$, etc. In general, we see recursively that $\Gamma(n) = (n-1)!$. (In 10B, one would prove this *by induction*.)

Another way to write the formula:

$$\Gamma(s+1) = s\Gamma(s).$$

There is a well known duplication formula for the Gamma function that is usually ascribed to Gauss. According to Wikipedia, it is also known as the *Legendre duplication formula*:

$$\Gamma(s)\Gamma(s+\frac{1}{2}) = 2^{1-2s}\sqrt{\pi}\,\Gamma(2s).$$

Take $s = \frac{1}{2}$:

$$\Gamma(\frac{1}{2})\Gamma(1) = 2^{0}\sqrt{\pi}\,\Gamma(1).$$

Because $\Gamma(1) = 1$, we get

$$\Gamma(\frac{1}{2}) = \sqrt{\pi}.$$

This is good to know because the $\sqrt{\pi}$ in the formula for the standard Gaussian is often written $\Gamma(1/2)$ in the literature.

For another example, we can calculate $\Gamma(3/2)$ in two ways. By the duplication formula,

$$\Gamma(3/2)\Gamma(2) = 2^{1-3}\sqrt{\pi}\,\Gamma(3),$$

so

$$\Gamma(3/2) = \frac{1}{1} \cdot \frac{1}{4}\sqrt{\pi} \cdot 2 = \frac{\sqrt{\pi}}{2}.$$

We can also write

$$\Gamma(3/2) = \Gamma(\frac{1}{2} + 1) = \frac{1}{2}\Gamma(\frac{1}{2}) = \frac{\sqrt{\pi}}{2}.$$

In the Math 10B notes, there's a critical moment where you read: "…Recall from last semester what you learned about the Gamma function…."

Last semester, I was just barely able to get some of my students to admit that they had been exposed to the Gamma function in 10A. I'm hoping that you all will nod enthusiastically next semester when the Gamma function arises.

Spoiler: you'll see Gamma values in the formula for the PDF of Chi Squared distributions, whatever they are. (You'll learn about them in 10B.)

In the Math 10B notes, there's a critical moment where you read: "…Recall from last semester what you learned about the Gamma function…."

Last semester, I was just barely able to get some of my students to admit that they had been exposed to the Gamma function in 10A. I'm hoping that you all will nod enthusiastically next semester when the Gamma function arises.

Spoiler: you'll see Gamma values in the formula for the PDF of Chi Squared distributions, whatever they are. (You'll learn about them in 10B.)

Student

"Student" was William Sealy Gosset, 1876–1937. He was the inventor of Student's *t*-test, which we will now explore.

The first thing to understand is that there are multiple distributions named after Student; one speaks of "Student's *t*-distribution," but there is a parameter in the distribution, the number of *degrees of freedom*. This number is ν ("nu"). Because ν can vary ($\nu = 1, 2, 3 \ldots$), there is really a family of distributions.

The number ν will be one less than the number of samples (*N* or *n*, depending on notation) that one draws.

Wikipedia:

> *Whereas a normal distribution describes a full population, t-distributions describe samples drawn from a full population; accordingly, the t-distribution for each sample size is different, and the larger the sample, the more the distribution resembles a normal distribution.*

What is t? It's a statistic like Z that's defined by a formula that's nearly a clone of the formula

$$Z = \frac{(\overline{X} - \mu)\sqrt{N}}{\sigma}.$$

The main difference is that t is applied in situations where μ is known but σ is unknown and is replaced by an estimate. Using the estimate in place of the (unkown) true value changes the distribution....

On Thursday we will see how the t-distributions are used for hypothesis testing.

Today we are just trying to figure out what the distributions *are*.

The νth Student PDF is given by

$$C \cdot (1 + \frac{x^2}{\nu})^{-(\nu+1)/2},$$

where C is a constant chosen so that the total area under the curve is 1 (as it's supposed to be).

If you graph $y = (1 + \frac{x^2}{\nu})^{-(\nu+1)/2}$ for different values of ν, you'll see a function that's symmetrical about the y-axis (because of the x^2) and that looks qualitatively like the PDF of the standard normal.

As $\nu \to \infty$, you can check (and please do check!) that $(1 + \frac{x^2}{\nu})^{-(\nu+1)/2} \to e^{-x^2/2}$.

To perform the check, you need to know the formula

$$\lim_{n\to\infty}\left(1+\frac{a}{n}\right)^n = e^a$$

when a is a real number. This can be verified (for example) by taking logs and using l'Hôpital's Rule.

Graphs of $(1+\frac{x^2}{\nu})^{-(\nu+1)/2}$ for $\nu = 1, 2, 3$:

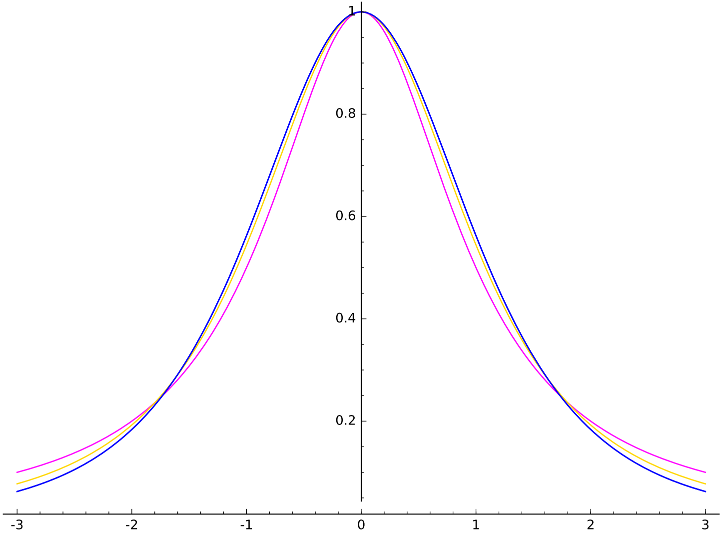

Color code: $\nu = 1$, magenta; $\nu = 2$, gold; $\nu = 3$, blue.

The constant C

We need to take

$$C = 1 \Big/ \int_{-\infty}^{\infty} (1 + \frac{x^2}{\nu})^{-(\nu+1)/2}\ dx.$$

The easiest case is that where $\nu = 1$. Because the antiderivative of $\dfrac{1}{1+x^2}$ is $\tan^{-1}(x)$,

$$\int_{-\infty}^{\infty} \frac{dx}{1+x^2} = \pi, \text{ so } C = 1/\pi.$$

In general,

$$C = \frac{1}{\sqrt{\nu\pi}} \frac{\Gamma(\frac{\nu+1}{2})}{\Gamma(\frac{\nu}{2})}.$$

This work when $\nu = 1$ because

$$\Gamma(\frac{\nu+1}{2}) = \Gamma(1) = 1, \quad \Gamma(\frac{\nu}{2}) = \Gamma(\frac{1}{2}) = \sqrt{\pi}.$$

According to Wikipedia, C involves π when ν is odd but not when ν is even.

About the final

It's cumulative but probably will have more on the last third of the course than on the second third, and more on the second third than on the first third. That's only because what's fresh in one's mind usually gets extra attention.

The final exam is on Friday night, 7–10PM, exam group 20.

However, we don't yet know where it is. *Do not show up in 2050 VLSB without first checking the final exam schedule!!!!*