This is the end

Kenneth A. Ribet

Math 10A
December 1, 2016

Plan for today

1. Announcements
2. Why it's true that $\Gamma(\frac{1}{2}) = \sqrt{\pi}$
3. Student t stuff
4. Posing for class photo
5. Course evaluations (after I leave)

Send email to sign up for the December 5 breakfast at 9AM.

Blue Bottle Coffee Company tomorrow at 9AM. Just show up.

Pop-in lunch on Friday, December 2 at 12:10PM. Just show up.

On Tuesday, we discussed the Gamma function

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

Using integration by parts, we saw that $\Gamma(n) = (n-1)!$ for integers $n \geq 1$.

What about the values of Γ at half-integers: $\Gamma(\frac{1}{2})$, $\Gamma(\frac{3}{2})$, $\Gamma(\frac{5}{2})$, etc.? Using the formula $\Gamma(s+1) = s\Gamma(s)$, we can compute these values once we know $\Gamma(\frac{1}{2})$. On Tuesday, I displayed the formula

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

and attributed it to Gauss. I said that the proof was hard.

I was wrong.

We have

$$\Gamma(\frac{1}{2}) = \int_0^\infty e^{-t} \frac{dt}{\sqrt{t}}$$

by definition. Let $x = \sqrt{t}$, $dx = \frac{1}{2} \frac{dt}{\sqrt{t}}$. The integral becomes

$$2 \int_0^\infty e^{-x^2} \, dx = \int_{-\infty}^\infty e^{-x^2} \, dx = \sqrt{\pi}.$$

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 ($\nu = 1, 2, 3 \ldots$), it is more appropriate to speak of a family of distributions.

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

The νth Student PDF is given by

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

where C is chosen (as usual) to make the total area under the curve equal to 1.

Encore presentation of graphs of $(1 + \frac{x^2}{\nu})^{-(\nu+1)/2}$:

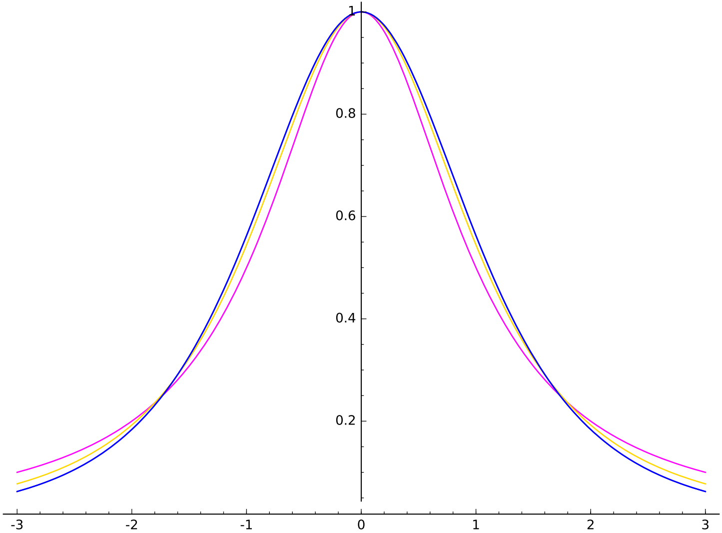

Color code: $\nu = 1$, magenta; $\nu = 2$, gold; $\nu = 3$, blue.

The graphs are symmetrical about the y-axis—this is because of the x^2 in the formula. (The functions are even.) Although the graphs look qualitatively like the PDF of the standard normal, their tails are much fatter than the tails of the normal PDF (at least for νs that are not too big).

As $\nu \to \infty$, you can verify that $(1 + \frac{x^2}{\nu})^{-(\nu+1)/2} \to e^{-x^2/2}$. This is a good exercise (or exam problem).

To do the exercise, you might start from 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.

As $\nu \to \infty$, you can verify that $(1 + \dfrac{x^2}{\nu})^{-(\nu+1)/2} \to e^{-x^2/2}$. This is a good exercise (or exam problem).

To do the exercise, you might start from 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.

The constant C

To get a PDF, we must take

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

The integral is not too hard to calculate when $\nu = 1$. 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 this case ($\nu = 1$).

In the general case, it's possible to write a compact formula for C using Γ-values:

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

This is compatible with our calculation for $\nu = 1$ because

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

Using the formulas that we've proved, you can check quickly that C involves π when ν is odd but not when ν is even. (See also Wikipedia.)

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 (unknown) true value changes the distribution....

It's best to consider an example, like the one from the `Prob-Stat` notes. This is a liberal rephrasing:

The average height of a college-age woman in the US is 64.7 inches. For some reason, we don't know the standard deviation of the height. We choose 10 UC women at random and find that their average height is 63.8 inches. Is this statistically significant?

The null hypothesis is that UC woman undergrads have an average height of 64.7 inches. The alternative hypothesis is that they don't. The analogue of the Z-statistic is

$$\frac{63.8 - 64.7}{\hat{\sigma}}\sqrt{10},$$

where $\hat{\sigma}$ is our estimate for the standard deviation (= the quantity that is withheld from us in this artificial example).

There are two wrinkles. The first is that we have to estimate the standard deviation from the heights of the 10 women who were lucky enough to be selected for this example. Call these heights $x_1, \ldots, x_{10}$. Their mean $\dfrac{x_1 + \cdots + x_{10}}{10}$ is given to us as 63.8. The estimate for the standard deviation that we use is

$$\hat{\sigma} = \left(\frac{1}{10-1} \sum_{i=1}^{10} (x_i - 63.8)^2 \right)^{1/2}.$$

This is called the sample standard deviation.

Why do we divide by $10 - 1$ instead of by 10? Why do we divide by $N - 1$ instead of by N? This is not a topic for the days of long shadows. There's an explanation in 10B.

The `prob-stat` notes give the value 16.2 for $\hat{\sigma}$. Using it, we find the value

$$t = \frac{63.8 - 64.7}{\hat{\sigma}}\sqrt{10} = -0.9\frac{\sqrt{10}}{\sqrt{16.2}} = -0.71.$$

We have to find the probability associated with "-0.71 or less" in the Student t-distribution with $10 - 1 = 9$ degrees of freedom.

For this we turn to the stat department online t-statistic calculator.

The probability of having a t-statistic between $-\infty$ and 0 is a huge 24%. The probability of being 0.71 or more standard deviations from the mean in this case is 48% (the double).

There is no way that we can reject the null hypothesis.

UC women are not short.

The probability of having a *t*-statistic between $-\infty$ and 0 is a huge 24%. The probability of being 0.71 or more standard deviations from the mean in this case is 48% (the double).

There is no way that we can reject the null hypothesis.

UC women are not short.

"A research study measured the pulse rates of 57 college men and found a mean pulse rate of 70.4211 beats per minute with a standard deviation of 9.9480 beats per minute. Researchers want to know if the mean pulse rate for all college men is different from the current standard of 72 beats per minute."

This is like the previous example; now

$$t = (70.4211 - 72) \cdot \frac{\sqrt{57}}{9.9480} \approx -1.20.$$

Two comments:

- The problem uses the phrase "standard deviation" rather than "sample standard deviation." Did they divide by 56 or by 57? There's no way for us to know.
- The mean and standard deviation are given to a ridiculous number of decimal places (four of them). There's some weird idea that lots of precision means that you're being very scientific.

We go back to the online calculator and change the number of degrees of freedom to 56. Between -1.2 and $+1.2$ is 76.48% of the area. Outside of this region is a bit more than 23.5% of the area. That's yuge.

There is no way that we reject the null hypothesis. The null hypothesis, by the way, is that the mean pulse rate of college men is the same as the vanilla mean pulse rate of 72.

My source for this example was https://onlinecourses.science.psu.edu/stat200/node/223. They get into a lot of trouble because they have no applet and have to consult pre-computed tables. They have pre-computed values for 50 and for 60 degrees of freedom and have to interpolate to figure out what happens for 56.

Modern devices are good.

Final remarks I

The final exam is 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.

The class will be split between 145 Dwinelle and the Pauley Ballroom. We will announce the set of GSIs whose sections go to Dwinelle; the other sectionds will go to Pauley.

Do not come to 2050 VLSB for the exam!!!!

Final remarks II

> *If you want to get to know him, you can, and even if you don't, you will.*

It has been a great pleasure for me to share Math 10A with you guys. If you want to keep in touch with me, it'll be pretty easy. We can continue breakfasts and pop-in lunches next semester. To organize them, we can use `piazza` (the class site should still be active) and Facebook. There is also email, and you can stroll by my office (885 Evans).

Course evaluations

The math department has recommended that we leave time at the end of our last lectures so that students can do course evaluations while in the classroom. You can can access the evaluation system through the following link:
https://course-evaluations.berkeley.edu.
As of early this morning, 31.28% of the class had completed their 10A evaluations.