Math 3081 (Probability and Statistics)

Lecture #23 of 27 $\sim$ August 12th, 2021

One-Sample t Tests

- One-Sample t Tests
- Foundation for Two-Sample t Tests

This material represents §4.2.3 from the course notes, and problems 1-5 from WeBWorK 7.

Recall, I

As mentioned last lecture, the t distribution arises in the situation where we want to do hypothesis tests on the sample mean from a normal distribution with unknown mean and unknown standard deviation.

- Suppose we draw $x_1, \ldots, x_n$ from a normal distribution with unknown mean μ and unknown standard deviation σ and let $\overline{x} = \frac{1}{n}(x_1 + \cdots + x_n)$ be the sample mean.
- We look at the normalized ratio $\dfrac{\overline{x} - \mu}{S/\sqrt{n}}$, which is the t-distribution analogue of the z score.

Recall, II

Recall also the pdf of the t distribution:

Definition

The <u>*t distribution with k degrees of freedom*</u> *is the continuous random variable T_k whose probability density function*

$$p_{T_k}(x) = \frac{\Gamma(\frac{k+1}{2})}{\sqrt{k\pi}\Gamma(\frac{k}{2})} \cdot (1 + x^2/k)^{-(k+1)/2}$$ *for all real numbers x.*

Recall, III

Our main result is the following:

Theorem (PDF of the t Distribution)

Suppose $n \geq 2$ and that $X_1, X_2, \dots, X_n$ are independent, identically normally distributed random variables with mean μ and standard deviation σ. If $\overline{X} = \frac{1}{n}(X_1 + \cdots + X_n)$ denotes the sample mean and $S = \sqrt{\frac{1}{n-1}\left[(X_1 - \overline{X})^2 + (X_2 - \overline{X})^2 + \cdots + (X_n - \overline{X})^2\right]}$ denotes the sample standard deviation, then the distribution of the normalized test statistic $\dfrac{\overline{X} - \mu}{S/\sqrt{n}}$ is the t distribution T_{n-1} with $n-1$ degrees of freedom.

Note that the number of degrees of freedom is $n-1$, one less than the number of data points.

One-Sample t Tests, I

Last lecture, we adapted the procedure for constructing confidence intervals with the normal distribution to construct confidence intervals using t statistics.

- Today, we adapt our procedures for z tests to do hypothesis testing with the t distribution: we call these t tests.
- We first describe one-sample t tests, in which we want to perform a hypothesis test on the unknown mean of a normal distribution with unknown standard deviation, based on an independent sampling of the distribution yielding n values $x_1, x_2, \ldots, x_n$.

One-Sample t Tests, II

The key difference between t tests and z tests is that the standard deviation of the normal distribution is unknown, rather than given to us.

- As usual with hypothesis tests, we first select appropriate null and alternative hypotheses and a significance level α.
- Our null hypothesis will be of the form H_0: $\mu = c$ for some constant c that is our hypothesized value for the mean of the normal distribution.
- We take the test statistic $t = \dfrac{\overline{x} - \mu}{S/\sqrt{n}}$, where $\overline{x}$ is the sample mean and S is the sample standard deviation.
- From our results about the t distribution, the distribution of the test statistic will be the t distribution T_{n-1} with $n-1$ degrees of freedom.

One-Sample t Tests, III

We can then calculate the p-value based on the alternative hypothesis.

- If the hypotheses are $H_0 : \mu = c$ and $H_a : \mu > c$, then the p-value is $P(T_{n-1} \geq t)$.
- If the hypotheses are $H_0 : \mu = c$ and $H_a : \mu < c$, then the p-value is $P(T_{n-1} \leq t)$.
- If the hypotheses are $H_0 : \mu = c$ and $H_a : \mu \neq c$, then the p-value is $P(|T_{n-1}| \geq |t|) = \begin{cases} 2P(T_{n-1} \geq t) & \text{if } t \geq \mu \\ 2P(T_{n-1} \leq t) & \text{if } t < \mu \end{cases}$.

We then compare the p-value to the significance level and then either reject or fail to reject the null hypothesis, as usual.

One-Sample t Tests, IV

Example: Suppose four values 9, 18, 7, 10 are sampled from a normal distribution with unknown mean and standard deviation. Test at the 20%, 11%, 2%, and 0.4% significance levels each hypothesis:

1. The mean is greater than 10.
2. The mean is greater than 0.
3. The mean is less than 25.
4. The mean is less than 5.
5. The mean is equal to 10.
6. The mean is equal to 16.

One-Sample t Tests, IV

Example: Suppose four values 9, 18, 7, 10 are sampled from a normal distribution with unknown mean and standard deviation. Test at the 20%, 11%, 2%, and 0.4% significance levels each hypothesis:

1. The mean is greater than 10.
2. The mean is greater than 0.
3. The mean is less than 25.
4. The mean is less than 5.
5. The mean is equal to 10.
6. The mean is equal to 16.

- First, we compute the sample mean $\hat{\mu} = \frac{1}{4}(9 + 18 + 7 + 10) = 11$ and sample standard deviation $S = \sqrt{\frac{1}{3}\left[(9-11)^2 + (18-11)^2 + (7-11)^2 + (10-11)^2\right]} = 4.8305$.

One-Sample t Tests, V

Example: Suppose four values 9, 18, 7, 10 are sampled from a normal distribution with unknown mean and standard deviation. Test at the 20%, 11%, 2%, and 0.4% significance levels each hypothesis:

1. The mean is greater than 10.

One-Sample t Tests, V

Example: Suppose four values 9, 18, 7, 10 are sampled from a normal distribution with unknown mean and standard deviation. Test at the 20%, 11%, 2%, and 0.4% significance levels each hypothesis:

1. The mean is greater than 10.

- Our hypotheses are $H_0 : \mu = 10$, $H_a : \mu > 10$; we want this one-sided alternative hypothesis since the actual sample mean is greater than 10.
- The value of our test statistic is $t = \dfrac{11 - 10}{4.8305/\sqrt{4}} = 0.4140$, giving p-value $P(T_{n-1} \geq 0.4140) = 0.3533$.
- Since this is greater than all four significance levels, we fail to reject the null hypothesis in all cases.

One-Sample t Tests, VI

Example: Suppose four values 9, 18, 7, 10 are sampled from a normal distribution with unknown mean and standard deviation. Test at the 20%, 11%, 2%, and 0.4% significance levels each hypothesis:

2. The mean is greater than 0.

One-Sample t Tests, VI

Example: Suppose four values 9, 18, 7, 10 are sampled from a normal distribution with unknown mean and standard deviation. Test at the 20%, 11%, 2%, and 0.4% significance levels each hypothesis:

2. The mean is greater than 0.
   - Our hypotheses are $H_0 : \mu = 0$, $H_a : \mu > 0$; we want this one-sided alternative hypothesis since the actual sample mean is greater than 0.
   - The value of our test statistic is $t = \dfrac{11 - 0}{4.8305/\sqrt{4}} = 4.5544$, giving p-value $P(T_{n-1} \geq 4.5544) = 0.00992$.
   - Since this is less than the first three significance levels, we reject the null hypothesis in those cases. However, it is greater than 0.4%, so we fail to reject the null hypothesis at that significance level.

One-Sample t Tests, VII

Example: Suppose four values 9, 18, 7, 10 are sampled from a normal distribution with unknown mean and standard deviation. Test at the 20%, 11%, 2%, and 0.4% significance levels each hypothesis:

3. The mean is less than 25.

One-Sample t Tests, VII

Example: Suppose four values 9, 18, 7, 10 are sampled from a normal distribution with unknown mean and standard deviation. Test at the 20%, 11%, 2%, and 0.4% significance levels each hypothesis:

3. The mean is less than 25.

- Our hypotheses are $H_0 : \mu = 25$, $H_a : \mu < 25$; we want this one-sided alternative hypothesis since the actual sample mean is less than 25.
- The value of our test statistic is $t = \dfrac{11 - 25}{4.8305/\sqrt{4}} = -5.7966$, giving p-value $P(T_{n-1} \leq -5.7966) = 0.00510$.
- Since this is less than the first three significance levels, we reject the null hypothesis in those cases. But it is greater than 0.4%, so we fail to reject in that case.

One-Sample t Tests, VIII

Example: Suppose four values 9, 18, 7, 10 are sampled from a normal distribution with unknown mean and standard deviation. Test at the 20%, 11%, 2%, and 0.4% significance levels each hypothesis:

4. The mean is less than 5.

One-Sample t Tests, VIII

Example: Suppose four values 9, 18, 7, 10 are sampled from a normal distribution with unknown mean and standard deviation. Test at the 20%, 11%, 2%, and 0.4% significance levels each hypothesis:

4. The mean is less than 5.

- Our hypotheses are $H_0 : \mu = 5$, $H_a : \mu > 5$; we want this one-sided alternative hypothesis since the actual sample mean is greater than 5.
- The value of our test statistic is $t = \dfrac{11 - 5}{4.8305/\sqrt{4}} = 2.4842$, giving p-value $P(T_{n-1} \geq 2.4842) = 0.0445$.
- Since this is less than 20% and 11%, we reject the null hypothesis in those cases. However, it is greater than 2% and 0.4%, so we fail to reject the null hypothesis at those significance levels.

One-Sample t Tests, IX

Example: Suppose four values 9, 18, 7, 10 are sampled from a normal distribution with unknown mean and standard deviation. Test at the 20%, 11%, 2%, and 0.4% significance levels each hypothesis:

5. The mean is equal to 10.

One-Sample t Tests, IX

Example: Suppose four values 9, 18, 7, 10 are sampled from a normal distribution with unknown mean and standard deviation. Test at the 20%, 11%, 2%, and 0.4% significance levels each hypothesis:

5. The mean is equal to 10.

- Our hypotheses are $H_0 : \mu = 0$, $H_a : \mu \neq 10$; we want this two-sided alternative hypothesis since we are only testing whether the mean equals 10 or not.
- The value of our test statistic is $t = \dfrac{11 - 10}{4.8305/\sqrt{4}} = 0.4140$, giving p-value $P(|T_{n-1}| \geq 0.4140) = 2P(T_{n-1} \geq 0.4140) = 0.7067$.
- Since this is (much!) greater than all of the listed significance levels, we fail to reject the null hypothesis in each case.

One-Sample t Tests, X

Example: Suppose four values 9, 18, 7, 10 are sampled from a normal distribution with unknown mean and standard deviation. Test at the 20%, 11%, 2%, and 0.4% significance levels each hypothesis:

6. The mean is equal to 16.

One-Sample t Tests, X

Example: Suppose four values 9, 18, 7, 10 are sampled from a normal distribution with unknown mean and standard deviation. Test at the 20%, 11%, 2%, and 0.4% significance levels each hypothesis:

6. The mean is equal to 16.

- Our hypotheses are $H_0 : \mu = 0$, $H_a : \mu \neq 16$; we want this two-sided alternative hypothesis since we are only testing whether the mean equals 16 or not.
- The value of our test statistic is $t = \dfrac{11 - 16}{4.8305/\sqrt{4}} = -2.0702$, giving p-value $P(|T_{n-1}| \geq |-2.0702|) = 2P(T_{n-1} \geq 2.0702) = 0.1302$.
- Since this is less than 20%, we reject the null hypothesis in that case. However, it is greater than 11%, 2% and 0.4%, so we fail to reject the null hypothesis at those significance levels.

One-Sample t Tests, XI

Example: To estimate the reaction yield, a new chemical synthesis is run three times, giving yields of 41.3%, 52.6%, and 56.1%. It is expected that the yield should be approximately normally distributed. Test at the 20%, 8%, and 1% significance levels the following hypotheses:

1. The average yield is above 45%.
2. The average yield is below 57%.
3. The average yield is above 64%.

One-Sample t Tests, XI

Example: To estimate the reaction yield, a new chemical synthesis is run three times, giving yields of 41.3%, 52.6%, and 56.1%. It is expected that the yield should be approximately normally distributed. Test at the 20%, 8%, and 1% significance levels the following hypotheses:

1. The average yield is above 45%.
2. The average yield is below 57%.
3. The average yield is above 64%.

- First, we compute the sample average $\hat{\mu} = \frac{1}{3}(41.3\% + 52.6\% + 56.1\%) = 50\%$, and the sample standard deviation $S = \sqrt{\frac{1}{2}\left[(41.3\% - 50\%)^2 + (52.6\% - 50\%)^2 + (56.1\% - 50\%)^2\right]} = 7.7350\%$.

One-Sample t Tests, XII

Example: To estimate the reaction yield, a new chemical synthesis is run three times, giving yields of 41.3%, 52.6%, and 56.1%. It is expected that the yield should be approximately normally distributed. Test at the 20%, 8%, and 1% significance levels the following hypotheses:

1. The average yield is above 45%.

One-Sample t Tests, XII

Example: To estimate the reaction yield, a new chemical synthesis is run three times, giving yields of 41.3%, 52.6%, and 56.1%. It is expected that the yield should be approximately normally distributed. Test at the 20%, 8%, and 1% significance levels the following hypotheses:

1. The average yield is above 45%.

- Our hypotheses are $H_0 : \mu = 45\%$, $H_a : \mu > 45\%$; we want this one-sided alternative hypothesis since the actual sample mean is greater than 45%.
- The value of our test statistic is $t = \dfrac{50\% - 45\%}{7.7350\%/\sqrt{3}} = 1.1196$, giving p-value $P(T_{n-1} \geq 1.1196) = 0.1896$.
- Since this is less than the first significance level 20%, we reject the null hypothesis in that case. However, it is greater than 8% and 1%, so we fail to reject the null hypothesis at those significance levels.

One-Sample t Tests, XIII

Example: To estimate the reaction yield, a new chemical synthesis is run three times, giving yields of 41.3%, 52.6%, and 56.1%. It is expected that the yield should be approximately normally distributed. Test at the 20%, 8%, and 1% significance levels the following hypotheses:

2. The average yield is below 57%.

One-Sample t Tests, XIII

Example: To estimate the reaction yield, a new chemical synthesis is run three times, giving yields of 41.3%, 52.6%, and 56.1%. It is expected that the yield should be approximately normally distributed. Test at the 20%, 8%, and 1% significance levels the following hypotheses:

2. The average yield is below 57%.

- Our hypotheses are $H_0 : \mu = 57\%$, $H_a : \mu < 57\%$; we want this one-sided alternative hypothesis since the actual sample mean is less than 57%.
- The value of our test statistic is $t = \dfrac{50\% - 57\%}{7.7350\%/\sqrt{3}} = -1.5675$, giving p-value $P(T_{n-1} \leq -1.5675) = 0.1288$.
- Since this is less than the first significance level 20%, we reject the null hypothesis in that case. However, it is greater than 8% and 1%, so we fail to reject the null hypothesis at those significance levels.

One-Sample t Tests, XIV

Example: To estimate the reaction yield, a new chemical synthesis is run three times, giving yields of 41.3%, 52.6%, and 56.1%. It is expected that the yield should be approximately normally distributed. Test at the 20%, 8%, and 1% significance levels the following hypotheses:

3. The average yield is above 64%.

One-Sample t Tests, XIV

<u>Example</u>: To estimate the reaction yield, a new chemical synthesis is run three times, giving yields of 41.3%, 52.6%, and 56.1%. It is expected that the yield should be approximately normally distributed. Test at the 20%, 8%, and 1% significance levels the following hypotheses:

3. The average yield is above 64%.

- Our hypotheses are $H_0 : \mu = 64\%$, $H_a : \mu < 64\%$; we want this one-sided alternative hypothesis since the actual sample mean is less than 64%.
- The value of our test statistic is $t = \dfrac{50\% - 64\%}{7.7350\%/\sqrt{3}} = -3.1349$, giving p-value $P(T_{n-1} \leq -3.1349) = 0.04423$.
- Since this is less than the first two significance levels 20% and 8%, we reject the null hypothesis in those cases. However, it is greater than 1%, so we fail to reject the null hypothesis at that significance level.

One-Sample t Tests, XV

Example: Eight recent polls of a candidate's approval rating yield results of 40%, 43%, 39%, 40%, 48%, 44%, 40%, and 38%. Assuming the polls are independent, test at the 11%, 3%, and 0.02% significance levels the hypotheses that

1. The candidate's approval rating is above 50%.
2. The candidate's approval rating is below 45%.
3. The candidate's approval rating is 40%.

One-Sample t Tests, XV

Example: Eight recent polls of a candidate's approval rating yield results of 40%, 43%, 39%, 40%, 48%, 44%, 40%, and 38%. Assuming the polls are independent, test at the 11%, 3%, and 0.02% significance levels the hypotheses that

1. The candidate's approval rating is above 50%.
2. The candidate's approval rating is below 45%.
3. The candidate's approval rating is 40%.

- First, we compute the sample average $\hat{\mu} = \frac{1}{8}(40\% + 43\% + 39\% + 40\% + 48\% + 44\% + 40\% + 38\%) = 41.5\%$, and the sample standard deviation $S = \sqrt{\frac{1}{7}\left[(40\% - 41.5\%)^2 + \cdots + (38\% - 41.5\%)^2\right]} = 3.2950\%$.

One-Sample t Tests, XV

Example: Eight recent polls of a candidate's approval rating yield results of 40%, 43%, 39%, 40%, 48%, 44%, 40%, and 38%. Test at the 11%, 3%, 0.4%, and 0.02% significance levels the hypotheses that

1. The candidate's approval rating is above 50%.

One-Sample t Tests, XV

Example: Eight recent polls of a candidate's approval rating yield results of 40%, 43%, 39%, 40%, 48%, 44%, 40%, and 38%. Test at the 11%, 3%, 0.4%, and 0.02% significance levels the hypotheses that

1. The candidate's approval rating is above 50%.

- Our hypotheses are $H_0 : \mu = 50\%$, $H_a : \mu < 50\%$; we want this one-sided alternative hypothesis since the actual sample mean is less than 50%.
- The value of our test statistic is $t = \dfrac{41.5\% - 50\%}{3.2950\%/\sqrt{8}} = -7.2964$, giving p-value $P(T_{n-1} \leq -7.2964) = 0.0000816$.
- Since this is less than all of the indicated significance levels, we reject the null hypothesis in all cases.

One-Sample t Tests, XVI

Example: Eight recent polls of a candidate's approval rating yield results of 40%, 43%, 39%, 40%, 48%, 44%, 40%, and 38%. Test at the 11%, 3%, 0.4%, and 0.02% significance levels the hypotheses that

2. The candidate's approval rating is below 45%.

One-Sample t Tests, XVI

Example: Eight recent polls of a candidate's approval rating yield results of 40%, 43%, 39%, 40%, 48%, 44%, 40%, and 38%. Test at the 11%, 3%, 0.4%, and 0.02% significance levels the hypotheses that

2. The candidate's approval rating is below 45%.

- Our hypotheses are $H_0 : \mu = 45\%$, $H_a : \mu < 45\%$; we want this one-sided alternative hypothesis since the actual sample mean is less than 45%.
- The value of our test statistic is $t = \dfrac{41.5\% - 45\%}{3.2950\%/\sqrt{8}} = -3.0044$, giving p-value $P(T_{n-1} \leq -3.0044) = 0.00991$.
- This is less than the first 2 significance levels so we reject the null hypothesis in those cases, but it is not less than the last 2, so we fail to reject there.

One-Sample t Tests, XVII

Example: Eight recent polls of a candidate's approval rating yield results of 40%, 43%, 39%, 40%, 48%, 44%, 40%, and 38%. Test at the 11%, 3%, 0.4%, and 0.02% significance levels the hypotheses that

3. The candidate's approval rating is 40%.

One-Sample t Tests, XVII

Example: Eight recent polls of a candidate's approval rating yield results of 40%, 43%, 39%, 40%, 48%, 44%, 40%, and 38%. Test at the 11%, 3%, 0.4%, and 0.02% significance levels the hypotheses that

3. The candidate's approval rating is 40%.

- Our hypotheses are $H_0 : \mu = 40\%$, $H_a : \mu \neq 40\%$; we want this two-sided alternative hypothesis since we are only testing whether the average equals 40%.
- The value of our test statistic is $t = \dfrac{41.5\% - 40\%}{3.2950\%/\sqrt{8}} = 1.2876$, giving p-value $2P(T_{n-1} \geq 1.2876) = 0.2388$.
- This is greater than all four significance levels, so we fail to reject the null hypothesis in all cases.

Confidence Intervals and t Tests, I

Just as with z tests, we can also interpret one-sample t tests using confidence intervals. The idea is exactly the same as before, except the underlying distribution is now a t distribution rather than a normal distribution.

- Since we work with the normalized test statistic, we have to compare to the corresponding normalized confidence interval, which is $(-t_{\alpha/2,df}, t_{\alpha/2,df})$.
- For a two-sided alternative hypothesis, if we give a $100(1-\alpha)\%$ confidence interval around the mean of a distribution under the conditions of the null hypothesis, then we will reject the null hypothesis with significance level α precisely when the sample statistic lies outside the normalized confidence interval.

Confidence Intervals and t Tests, II

Here is the corresponding picture:

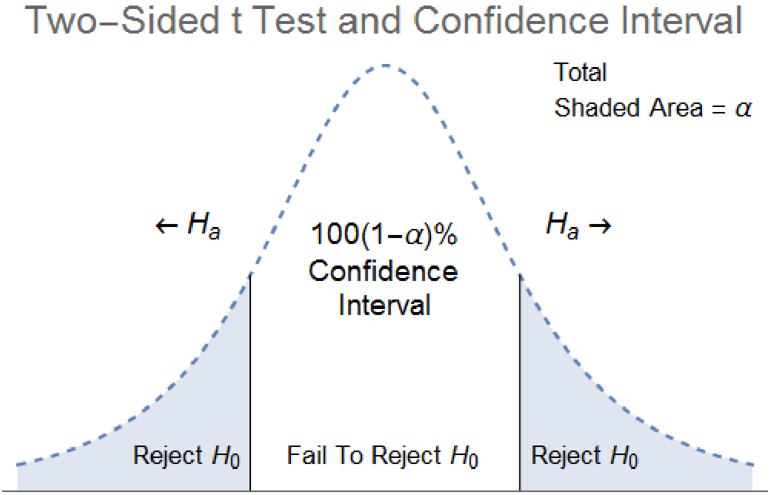

Confidence Intervals and t Tests, III

We can do the same thing with a one-sided alternative hypothesis, but because of the lack of symmetry in the rejection region, we instead need to use a $100(1-2\alpha)\%$ confidence interval to get the correct area.

Confidence Intervals and t Tests, IV

Here is the corresponding picture:

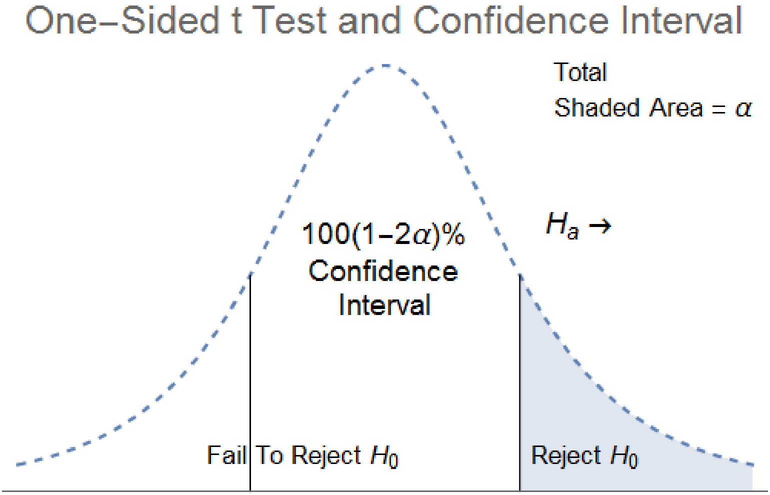

Confidence Intervals and t Tests, V

Example: The online list prices for four randomly-chosen statistics textbooks are \$193.95, \$171.89, \$221.80, and \$215.32. Assume statistics textbook prices are approximately normally distributed.

1. Find 80%, 90%, 95%, 98%, 99%, and 99.5% confidence intervals for the average list price of a statistics textbook.
2. Test at the 10% and 1% levels that the average price is \$200.
3. Test at the 10% and 1% levels that the average price is \$230.
4. Test at the 10% and 1% levels that the average price is \$275.
5. Test at the 10% and 1% levels that the average price is above \$170.
6. Test at the 10% and 1% levels that the average price is above \$270.

Confidence Intervals and t Tests, VI

Example: The online list prices for four randomly-chosen statistics textbooks are \$193.95, \$171.89, \$221.80, and \$215.32. Assume statistics textbook prices are approximately normally distributed.

1. Find 80%, 90%, 95%, 98%, 99%, and 99.5% confidence intervals for the average list price of a statistics textbook.

Confidence Intervals and t Tests, VI

Example: The online list prices for four randomly-chosen statistics textbooks are \$193.95, \$171.89, \$221.80, and \$215.32. Assume statistics textbook prices are approximately normally distributed.

1. Find 80%, 90%, 95%, 98%, 99%, and 99.5% confidence intervals for the average list price of a statistics textbook.

- The sample mean is $\hat{\mu} = \frac{1}{4}(\$193.95 + \$171.89 + \$221.80 + \$215.32) = \$200.74$ with sample standard deviation $S = \sqrt{\frac{1}{3}\left[(\$193.95 - \$200.74)^2 + \cdots + (\$215.32 - \$200.74)^2\right]} = \$22.617$.
- To find the confidence intervals, we need to look up or calculate the appropriate t-statistics for the given confidence levels and $n = 4$ (3 degrees of freedom).

Confidence Intervals and t Tests, VII

Example: The online list prices for four randomly-chosen statistics textbooks are \$193.95, \$171.89, \$221.80, and \$215.32. Assume statistics textbook prices are approximately normally distributed.

1. Find 80%, 90%, 95%, 98%, 99%, and 99.5% confidence intervals for the average list price of a statistics textbook.

The confidence intervals are as follows:

$1-\alpha$	Confidence Interval
80%	(\$182.22, \$219.26)
90%	(\$174.13, \$227.35)
95%	(\$164.75, \$236.73)
98%	(\$149.39, \$252.09)
99%	(\$134.69, \$266.79)
99.5%	(\$116.46, \$285.02)

Confidence Intervals and t Tests, VI

Example: The online list prices for four randomly-chosen statistics textbooks are \$193.95, \$171.89, \$221.80, and \$215.32. Assume statistics textbook prices are approximately normally distributed.

2. Test at the 10% and 1% levels that the average price is \$200.

Confidence Intervals and t Tests, VI

Example: The online list prices for four randomly-chosen statistics textbooks are \$193.95, \$171.89, \$221.80, and \$215.32. Assume statistics textbook prices are approximately normally distributed.

2. Test at the 10% and 1% levels that the average price is \$200.

- We take $H_0 : \mu = 200$, $H_a : \mu \neq 200$. This is a two-sided confidence interval, and so we want to look at the $100(1-\alpha)\%$ confidence intervals for $\alpha = 0.10$ and $\alpha = 0.01$.

$1-\alpha$	Confidence Interval
90%	$(\$174.13, \$227.35)$
99%	$(\$134.69, \$266.79)$

- Since 200 lies in both the 90% and 99% confidence intervals, we fail to reject the null hypothesis in both cases.
- Explicitly, the normalized test statistic is $t = 0.0654$, and so our p-value is $2P(T_{n-1} \geq 0.0654) = 0.9510$.

Confidence Intervals and t Tests, VI

Example: The online list prices for four randomly-chosen statistics textbooks are \$193.95, \$171.89, \$221.80, and \$215.32. Assume statistics textbook prices are approximately normally distributed.

3. Test at the 10% and 1% levels that the average price is \$230.

Confidence Intervals and t Tests, VI

Example: The online list prices for four randomly-chosen statistics textbooks are \$193.95, \$171.89, \$221.80, and \$215.32. Assume statistics textbook prices are approximately normally distributed.

3. Test at the 10% and 1% levels that the average price is \$230.

- We take $H_0 : \mu = 230$, $H_a : \mu \neq 230$. As before, we want to look at the 90% and 99% confidence intervals.

$1-\alpha$	Confidence Interval
90%	$(\$174.13, \$227.35)$
99%	$(\$134.69, \$266.79)$

- Since 230 lies outside the 90% confidence interval, we reject the null hypothesis at the 10% significance level.
- But since 230 lies inside the 99% confidence interval, we fail to reject the null hypothesis at the 1% significance level.
- The actual p-value is $2P(T_{n-1} \leq -2.5875) = 0.0812$: below 10% but above 1%.

Confidence Intervals and t Tests, VI

Example: The online list prices for four randomly-chosen statistics textbooks are \$193.95, \$171.89, \$221.80, and \$215.32. Assume statistics textbook prices are approximately normally distributed.

4. Test at the 10% and 1% levels that the average price is \$275.

Confidence Intervals and t Tests, VI

Example: The online list prices for four randomly-chosen statistics textbooks are \$193.95, \$171.89, \$221.80, and \$215.32. Assume statistics textbook prices are approximately normally distributed.

4. Test at the 10% and 1% levels that the average price is \$275.

- We take $H_0 : \mu = 275$, $H_a : \mu \neq 275$. As before, we want to look at the 90% and 99% confidence intervals.

$1 - \alpha$	Confidence Interval
90%	$(\$174.13, \$227.35)$
99%	$(\$134.69, \$266.79)$

- Since 275 lies outside both the 90% and 99% confidence intervals, we reject the null hypothesis in both cases.
- The p-value is $2P(T_{n-1} \leq -6.5669) = 0.00278$, below 1%.

Confidence Intervals and t Tests, VII

Example: The online list prices for four randomly-chosen statistics textbooks are \$193.95, \$171.89, \$221.80, and \$215.32. Assume statistics textbook prices are approximately normally distributed.

5. Test at the 10% and 1% levels that the avg price is $>$ \$170.

Confidence Intervals and t Tests, VII

Example: The online list prices for four randomly-chosen statistics textbooks are \$193.95, \$171.89, \$221.80, and \$215.32. Assume statistics textbook prices are approximately normally distributed.

5. Test at the 10% and 1% levels that the avg price is $>$ \$170.

- We take $H_0 : \mu = 170$, $H_a : \mu > 170$. Now we have a one-sided alternative hypothesis, so we want to look at the $100(1-2\alpha)\%$ confidence intervals.

α	Confidence Interval
80%	$(\$182.22, \$219.26)$
98%	$(\$149.39, \$252.09)$

- Since 170 lies below the 80% confidence interval, we reject the null hypothesis at the 10% significance level.
- However, 170 does lie inside the 98% confidence interval, so we fail to reject the null hypothesis at the 1% level.
- The p-value is $P(T_{n-1} \geq 2.7184) = 0.02653$.

Confidence Intervals and t Tests, VIII

Example: The online list prices for four randomly-chosen statistics textbooks are \$193.95, \$171.89, \$221.80, and \$215.32. Assume statistics textbook prices are approximately normally distributed.

6. Test at the 10% and 1% levels that the avg price is $>$ \$270.

Confidence Intervals and t Tests, VIII

Example: The online list prices for four randomly-chosen statistics textbooks are \$193.95, \$171.89, \$221.80, and \$215.32. Assume statistics textbook prices are approximately normally distributed.

6. Test at the 10% and 1% levels that the avg price is $>$ \$270.

- Based on the statement we try $H_0 : \mu = 270$, $H_a : \mu > 270$. As above, we want the $100(1 - 2\alpha)\%$ confidence intervals.

$1-\alpha$	Confidence Interval
80%	$(\$182.22, \$219.26)$
98%	$(\$149.39, \$252.09)$

- Notice that 270 lies outside both the 80% and 98% confidence intervals. However (!), the confidence intervals themselves are below 270, meaning that our deviation away from the hypothesized value falls into the null hypothesis tail of the distribution, rather than the alternative hypothesis tail.

Confidence Intervals and t Tests, IX

Example: The online list prices for four randomly-chosen statistics textbooks are \$193.95, \$171.89, \$221.80, and \$215.32. Assume statistics textbook prices are approximately normally distributed.

6. Test at the 10% and 1% levels that the avg price is $>$ \$270.

- Thus, we fail to reject the null hypothesis at either significance level.
- The actual p-value here is $P(T_{n-1} \geq -6.1247) = 0.9982$.

Confidence Intervals and t Tests, IX

Example: The online list prices for four randomly-chosen statistics textbooks are \$193.95, \$171.89, \$221.80, and \$215.32. Assume statistics textbook prices are approximately normally distributed.

6. Test at the 10% and 1% levels that the avg price is $>$ \$270.

- Thus, we fail to reject the null hypothesis at either significance level.
- The actual p-value here is $P(T_{n-1} \geq -6.1247) = 0.9982$.
- Here, we should have used the alternative hypothesis $H_a : \mu < 270$, since the sample mean was less than 270.
- In this case, 270 would still lie outside both the 90% and 99% confidence intervals, but now 270 would land in the alternative hypothesis tail rather than the null hypothesis tail, so we would reject the null hypothesis at both significance levels.
- In that situation, the p-value is $P(T_{n-1} \leq -6.1247) = 0.001800$, which is indeed quite small.

Two-Sample t Tests, I

Now that we have treated the situation of one-sample t tests, we will get into the thornier issue of two-sample t tests, in which we want to compare the unknown means of two normally-distributed populations with unknown standard deviations.

Two-Sample t Tests, II

Let's first review the setup for a two-sample z test (this also serves as good review!):

- Suppose the two populations are labeled A and B, with respective means μ_A and μ_B and population standard deviations σ_A and σ_B.
- We sample population A a total of n_A times, and population B a total of n_B times.
- We would like to take our test statistic as the difference $\mu_A - \mu_B$ in the two population means.
- First suppose that we are testing whether $\mu_A = \mu_B$, which we can equivalently phrase as asking whether $\mu_A - \mu_B = 0$.

Two-Sample t Tests, III

- The sample mean $\hat{\mu}_A$ will be normally distributed with mean μ_A and standard deviation $\sigma_A/\sqrt{n_A}$.
- Likewise, $\hat{\mu}_B$ will be normally distributed with mean μ_B and standard deviation $\sigma_B/\sqrt{n_B}$.
- The key piece of information here is that since $\hat{\mu}_A$ and $\hat{\mu}_B$ are independent and normally distributed, their difference is also normally distributed with mean $\mu_A - \mu_B$ and standard deviation $\sqrt{\frac{\sigma_A^2}{n_A} + \frac{\sigma_B^2}{n_B}}$.
- Therefore, if we are testing the null hypothesis H_0: $\mu_A - \mu_B = c$, then under the assumption of the null hypothesis, our test statistic $\hat{\mu}_A - \hat{\mu}_B$ will be normally distributed with mean c and standard deviation $\sqrt{\frac{\sigma_A^2}{n_A} + \frac{\sigma_B^2}{n_B}}$.

Two-Sample t Tests, IV

That's all well and good when we know the standard deviations of A and B. But now suppose we don't.

- Then we must use the sample standard deviation estimates S_A and S_B to estimate the standard deviation of $\hat{\mu}_A - \hat{\mu}_B$.
- However, just as we discussed before, using the sample standard deviation in place of the population standard deviation changes the underlying distributions.
- Although $\dfrac{\hat{\mu}_A - \mu_A}{\sigma_A/\sqrt{n_A}}$ has the standard normal distribution, $\dfrac{\hat{\mu}_A - \mu_A}{S_A/\sqrt{n_A}}$ has the distribution of the t-distributed random variable T_{n_A-1}.

Two-Sample t Tests, V

- If we solve for the distribution of $\hat{\mu}_A$, we see it is no longer given by the normal random variable $\mu_A + \frac{\sigma}{\sqrt{n_A}} N_{0,1}$ (normal with with mean μ_A and standard deviation $\sigma_A/\sqrt{n_A}$), but rather a "rescaled" t distribution $\mu_A + \frac{S_A}{\sqrt{n_A}} T_{n_A-1}$.
- Likewise, the random variable $\hat{\mu}_B$ has a rescaled t distribution $\mu_B + \frac{S_B}{\sqrt{n_B}} T_{n_B-1}$.
- Then the quantity $\hat{\mu}_A - \hat{\mu}_B$ is modeled by the random variable $\left[\mu_A + \frac{S_A}{\sqrt{n_A}} T_{n_A-1}\right] - \left[\mu_B + \frac{S_B}{\sqrt{n_B}} T_{n_B-1}\right] =$ $(\mu_A - \mu_B) + \frac{S_A}{\sqrt{n_A}} T_{n_A-1} - \frac{S_B}{\sqrt{n_B}} T_{n_B-1}$.

(Ugh!)

Two-Sample t Tests, VI

To summarize, all of those calculations say $\mu_A - \mu_B$ is modeled by $(\hat{\mu}_A - \hat{\mu}_B) + \frac{S_A}{\sqrt{n_A}} T_{n_A - 1} - \frac{S_B}{\sqrt{n_B}} T_{n_B - 1} \dots$ (whatever that is!)

- The problem is that we don't have a nice description of what the difference between two (scaled) t distributions looks like.
- For normal distributions, we can use the very convenient fact that the sum or difference of normal random variables is also normal; that is not the case for t distributions!
- In principle, because we know the probability density functions of $T_{n_A - 1}$ and $T_{n_B - 1}$, and they are independent, we could calculate the probability density function of the random variable listed above for particular values of all of the parameters.
- But that does not solve our problem, because we need to write down a distribution that is independent of the test parameters (i.e., that does not depend on S_A and S_B).

Two-Sample t Tests, VII

So... how do we set up two-sample t tests?

... I'll tell you next week. Happy studying!

Summary

We discussed one-sample t tests.

We discussed the relationship between one-sample t tests and confidence intervals.

We briefly motivated two-sample t tests.

Next lecture: Two-sample t tests.