Math 3081 (Probability and Statistics)

Lecture #16 of 27 $\sim$ August 2, 2021

Interval Estimation (Part 1)

- Estimation, Accuracy, and Precision
- Confidence Intervals
- Normal Confidence Intervals

This material represents §3.2.1-3.2.2 from the course notes, and problems 9-12 from WeBWorK 5.

Motivation, I

When estimating an unknown parameter, it is (of course) desirable to have a prediction that is as accurate as possible, but it is also important say how accurate we expect the prediction to be.

- For example, if we are estimating the height of a building, an estimate of 25.43 meters is certainly useful, but it is far more useful if we can also say that it is correct to within 0.01 meters.
- In contrast, if we estimate the height to be 25.43 meters but that our estimate is correct only to within 20 meters, the estimate is not nearly as good!

Motivation, II

This is related to the issue of measurement precision, but is not exactly the same.

- If we measure the height three times and obtain estimates of 25.43 meters, 25.41 meters, and 25.44 meters, we can be more confident in the overall accuracy than if the three measurements were 25.43 meters, 17.42 meters, and 33.15 meters.
- Nonetheless, these measurements by themselves do not provide an explicit error range for our estimated height.

Motivation, III

What we are seeking is to expand our discussion from pointwise parameter estimates, where we estimate the actual value of the parameter, to interval estimates, where we give an interval that we believe the parameter should lie in.

- Our discussion of unbiasedness is partially in this direction, since unbiasedness eliminates the existence of systematic error (i.e., error that tends to bias the estimate either too high or too low on average).
- Our discussion of estimator efficiency also represents partial progress toward this goal: efficient estimators have a smaller variance, and thus (by definition) will display less variation in their values than less-efficient estimators.

Motivation, IV

However, neither unbiasedness nor efficiency measures what we are looking for.

- Unbiasedness is only an average measure, and doesn't tell us anything about a specific measurement.
- Efficiency is a measurement of precision (the closeness of the measurements to one another) rather than of accuracy (the closeness of the measurements to the true value).
- What we are seeking is a way to quantify the accuracy of our estimations.

Confidence Intervals, I

One approach to quantifying the uncertainty in our estimates is to construct a confidence interval: this is an interval around our estimated value in which we believe the true value should lie.

- Of course, since the estimator is itself defined in terms of the values of a random sample, we cannot generally be completely certain that the true value of the parameter lies in any useful interval we could define. (We could, of course, simply declare our confidence interval to be the entire real line, but this would not give a useful prediction!)
- But what we can do is compute the probability that the true parameter value lies in the interval we give. If the probability is reasonably large (depending on the context, one may consider values such as 50%, or 90%, or 95%, or 99%, or 99.99% as appropriately large probabilities) then we can be reasonably confident in the accuracy of our estimation.

Confidence Intervals, II

Definition

If X is a random variable and $0 < \alpha < 1$, a $100(1-\alpha)\%$ confidence interval for X is an interval (a, b) with $a < X < b$ such that $P(a < X < b) = 1 - \alpha$.

- We use the notation $100(1-\alpha)\%$ is because it is traditional to quote the size of the confidence interval as a percent, rather than as a raw probability. Thus, for example, a 95% confidence interval for X is an interval (a, b) where X should land 95% of the time.
- In principle, one can define confidence intervals for any random variable, but in practice they are only given for random variables that represent parameter estimators.

Confidence Intervals, III

When θ is an unknown parameter, we interpret a confidence interval for θ as giving us a "reasonable error range" (for a precisely quantified notion of reasonable, determined by the error probability α) on a specific estimation $\hat{\theta}$ for θ that we have computed.

Confidence Intervals, IV

Example: Suppose we perform a maximum likelihood estimate for the parameter $\lambda = \theta$ of a Poisson distribution and obtain the estimate $\hat{\theta} = 1.39$, and by analysis of the variation of the estimator we are able to determine that there is a 95% probability that the true value of θ lies in the interval $(1.33, 1.51)$.

- This interval $(1.33, 1.51)$ is then a 95% confidence interval for our estimate, and it provides substantial additional context to our pointwise estimate $\hat{\theta} = 1.39$, since it quantifies how much variation we should expect to see in the true value of the parameter.
- If we sampled this distribution repeatedly and constructed a 95% confidence interval using each sample, we would expect the true value of the parameter to lie inside the interval 95% of the time.

Confidence Intervals, V

When we are constructing confidence intervals using parameter estimates, we typically will want to work with unbiased estimators that are as efficient as possible.

- If the estimator is unbiased, then the confidence interval will not tend to be biased above or below the true value of the parameter (i.e., it yields better average accuracy from a given data sample).
- If the estimator is efficient, then the size of the interval will be as small as possible, which yields tighter estimates for a given confidence level (i.e., it yields better overall precision from a given data sample).

Confidence Intervals, VI

In general, computing a confidence interval requires being able to analyze the precise nature of the variation in the estimator $\hat{\theta}$ relative to the true value θ.

In certain situations, we can describe this variation quite precisely, but in others it can be very difficult.

We will start by treating one of the simplest cases of computing a confidence interval: estimating the mean of a normal distribution whose standard deviation is known.

Normal Confidence Intervals, I

So suppose we sample a normal distribution with unknown mean μ and known standard deviation σ, obtaining values $x_1, x_2, \ldots, x_n$: our goal is to give a confidence interval for μ.

- We have previously shown that the maximum likelihood estimator for the mean, which is simply the sample mean $\hat{\mu} = \frac{1}{n}(x_1 + x_2 + \cdots + x_n)$, is unbiased and is the most efficient unbiased estimator for μ.
- Furthermore, from our results on the normal distribution and the central limit theorem, we know that since the x_i are independent and normally distributed with mean μ and standard deviation σ, the sample mean $\hat{\mu} = \frac{1}{n}(x_1 + x_2 + \cdots + x_n)$ will also be normally distributed with mean μ and standard deviation $\sigma/\sqrt{n}$.

Normal Confidence Intervals, II

So far, our analysis has proceeded as if we knew μ and wanted to understand the variation in $\hat{\mu}$.

- But now we can switch our focus from the variation of $\hat{\mu}$ given μ to the variation of μ given $\hat{\mu}$: from the previous slide, we know that the difference $\hat{\mu} - \mu$ is normally distributed with mean 0 and standard deviation $\sigma/\sqrt{n}$.
- This is the same as saying that the value of μ is normally distributed with mean $\hat{\mu}$ and standard deviation $\sigma/\sqrt{n}$.

Normal Confidence Intervals, II

So far, our analysis has proceeded as if we knew μ and wanted to understand the variation in $\hat{\mu}$.

- But now we can switch our focus from the variation of $\hat{\mu}$ given μ to the variation of μ given $\hat{\mu}$: from the previous slide, we know that the difference $\hat{\mu} - \mu$ is normally distributed with mean 0 and standard deviation $\sigma/\sqrt{n}$.
- This is the same as saying that the value of μ is normally distributed with mean $\hat{\mu}$ and standard deviation $\sigma/\sqrt{n}$.
- But now, because we have both $\hat{\mu}$ and $\sigma/\sqrt{n}$, to construct a confidence interval for μ we just need to compute the necessary probabilities using the normal distribution.
- Explicitly, if $N_{\hat{\mu},\sigma/\sqrt{n}}$ is the normal distribution with mean $\hat{\mu}$ and standard deviation $\sigma/\sqrt{n}$, then $P(a < \mu < b) = P(a < N_{\hat{\mu},\sigma/\sqrt{n}} < b)$.

Normal Confidence Intervals, III

We can therefore construct a $100(1-\alpha)\%$ confidence interval for μ simply by finding a range (a, b) such that $P(a < N_{\hat{\mu},\sigma/\sqrt{n}} < b) = 1 - \alpha$, as illustrated in the diagram below:

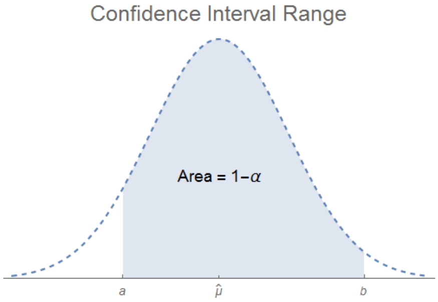

Normal Confidence Intervals, IV

There are many possible choices for this interval.

- To narrow things down, we usually require that the interval be symmetric around $\hat{\mu}$, which has the effect of making the interval as short as possible.
- For convenience we can also rephrase this condition in terms of the standard normal distribution $N_{0,1}$ by rescaling.
- Specifically, if we compute the constant c such that $P(-c < N_{0,1} < c) = 1 - \alpha$, then this yields the $100(1-\alpha)\%$ confidence interval $(a, b) = (\hat{\mu} - c\frac{\sigma}{\sqrt{n}}, \hat{\mu} + c\frac{\sigma}{\sqrt{n}})$.
- You can think of this calculation as "finding the number c of standard deviations away from the mean that will capture a total area of $1 - \alpha$".

Normal Confidence Intervals, V

We can also compute the value c for the confidence interval using the inverse cdf for the standard normal distribution.

- Specifically, since the two tails of the normal distribution have equal area, that means $P(N < -c) = P(N > c)$.
- Thus, if $P(-c < N_{0,1} < c) = 1 - \alpha$, then since $P(N < -c) + P(-c < N < c) + P(N > c) = 1$, that means $P(N_{0,1} < -c) = \alpha/2$, and thus $P(N_{0,1} < c) = (1 + \alpha)/2$.
- This allows us to compute the value of c by evaluating the inverse cumulative distribution function for $N_{0,1}$: specifically, c is the value of the inverse cdf on the value $(1 + \alpha)/2$.

Normal Confidence Intervals, VI

We can summarize the results of this discussion as follows:

Proposition (Normal Confidence Intervals)

A $100(1-\alpha)\%$ confidence interval for the unknown mean μ of a normal distribution with known standard deviation σ is given by $\hat{\mu} \pm c\dfrac{\sigma}{\sqrt{n}} = (\hat{\mu} - c\dfrac{\sigma}{\sqrt{n}}, \hat{\mu} + c\dfrac{\sigma}{\sqrt{n}})$ where n sample points $x_1, \dots, x_n$ are taken from the distribution, $\hat{\mu} = \frac{1}{n}(x_1 + \cdots + x_n)$ is the sample mean, and c is the constant satisfying $P(-c < N_{0,1} < c) = 1 - \alpha$.

Here are various pairs (α, c) where $P(-c < N_{0,1} < c) = 1 - \alpha$:

$1-\alpha$	50%	80%	90%	95%	98%	99%	99.5%	99.9%
c	0.6745	1.2816	1.6449	1.9600	2.3263	2.5758	2.8070	3.2905

Normal Confidence Intervals, VII

The term $c \frac{\sigma}{\sqrt{n}}$ is called the <u>margin of error</u> for the confidence interval, since it represents the maximum distance away (in either direction) values in the interval can be from the center.

- If we imagine choosing different sample sizes n, we can see that the margin of error in the estimate decreases with larger n. This is, of course, quite intuitive: if we sample more values, we would expect the errors to tend to cancel one another out on average, yielding an average that is more likely to land close to the true value than any single observation. (More formally, it follows from the central limit theorem.)
- More precisely, the margin of error is proportional to $1/\sqrt{n}$: so, for example, to cut the margin of error in half would require a sample size that is 4 times as large.

Examples, I

Example: A normal distribution with unknown mean μ and standard deviation $\sigma = 1$ is sampled four times, yielding the values 1.4, 0.2, 2.9, and 1.1.

1. Find a 50% confidence interval for μ.
2. Find a 90% confidence interval for μ.
3. Find a 95% confidence interval for μ.
4. Find a 99% confidence interval for μ.

Examples, I

Example: A normal distribution with unknown mean μ and standard deviation $\sigma = 1$ is sampled four times, yielding the values 1.4, 0.2, 2.9, and 1.1.

1. Find a 50% confidence interval for μ.
2. Find a 90% confidence interval for μ.
3. Find a 95% confidence interval for μ.
4. Find a 99% confidence interval for μ.

- We need only apply the formula and look up the proper value of c in the table:

$1-\alpha$	50%	80%	90%	95%	98%	99%	99.5%	99.9%
c	0.6745	1.2816	1.6449	1.9600	2.3263	2.5758	2.8070	3.2905

Examples, II

Example: A normal distribution with unknown mean μ and standard deviation $\sigma = 1$ is sampled four times, yielding the values 1.4, 0.2, 2.9, and 1.1.

$1-\alpha$	50%	80%	90%	95%	98%	99%	99.5%	99.9%
c	0.6745	1.2816	1.6449	1.9600	2.3263	2.5758	2.8070	3.2905

1. Find a 50% confidence interval for μ.

Examples, II

Example: A normal distribution with unknown mean μ and standard deviation $\sigma = 1$ is sampled four times, yielding the values 1.4, 0.2, 2.9, and 1.1.

$1-\alpha$	50%	80%	90%	95%	98%	99%	99.5%	99.9%
c	0.6745	1.2816	1.6449	1.9600	2.3263	2.5758	2.8070	3.2905

1. Find a 50% confidence interval for μ.

- We have $n = 4$, $\hat{\mu} = \frac{1}{4}(1.4 + 0.2 + 2.9 + 1.1) = 1.4$, $\sigma/\sqrt{n} = 0.5$.
- So the 50% confidence interval is $\hat{\mu} \pm 0.6745 \cdot \sigma/\sqrt{n} = (1.063, 1.737)$.

Examples, III

Example: A normal distribution with unknown mean μ and standard deviation $\sigma = 1$ is sampled four times, yielding the values 1.4, 0.2, 2.9, and 1.1.

$1-\alpha$	50%	80%	90%	95%	98%	99%	99.5%	99.9%
c	0.6745	1.2816	1.6449	1.9600	2.3263	2.5758	2.8070	3.2905

2. Find a 90% confidence interval for μ.

Examples, III

Example: A normal distribution with unknown mean μ and standard deviation $\sigma = 1$ is sampled four times, yielding the values 1.4, 0.2, 2.9, and 1.1.

$1-\alpha$	50%	80%	90%	95%	98%	99%	99.5%	99.9%
c	0.6745	1.2816	1.6449	1.9600	2.3263	2.5758	2.8070	3.2905

2. Find a 90% confidence interval for μ.
   - This is $\hat{\mu} \pm 1.6449 \cdot \sigma/\sqrt{n} = (0.577, 2.223)$.
3. Find a 95% confidence interval for μ.

Examples, III

<u>Example</u>: A normal distribution with unknown mean μ and standard deviation $\sigma = 1$ is sampled four times, yielding the values 1.4, 0.2, 2.9, and 1.1.

$1-\alpha$	50%	80%	90%	95%	98%	99%	99.5%	99.9%
c	0.6745	1.2816	1.6449	1.9600	2.3263	2.5758	2.8070	3.2905

2. Find a 90% confidence interval for μ.

- This is $\hat{\mu} \pm 1.6449 \cdot \sigma/\sqrt{n} = (0.577, 2.223)$.

3. Find a 95% confidence interval for μ.

- This is $\hat{\mu} \pm 1.9600 \cdot \sigma/\sqrt{n} = (0.420, 2.380)$.

4. Find a 99% confidence interval for μ.

Examples, III

Example: A normal distribution with unknown mean μ and standard deviation $\sigma = 1$ is sampled four times, yielding the values 1.4, 0.2, 2.9, and 1.1.

$1-\alpha$	50%	80%	90%	95%	98%	99%	99.5%	99.9%
c	0.6745	1.2816	1.6449	1.9600	2.3263	2.5758	2.8070	3.2905

2. Find a 90% confidence interval for μ.
   - This is $\hat{\mu} \pm 1.6449 \cdot \sigma/\sqrt{n} = (0.577, 2.223)$.
3. Find a 95% confidence interval for μ.
   - This is $\hat{\mu} \pm 1.9600 \cdot \sigma/\sqrt{n} = (0.420, 2.380)$.
4. Find a 99% confidence interval for μ.
   - This is $\hat{\mu} \pm 2.5758 \cdot \sigma/\sqrt{n} = (0.112, 2.688)$.

Examples, IV

Example: The diameters of bolts manufactured at Factory X are distributed normally with standard deviation 0.01mm. A random sample of 10 bolts from one lot has average diameter 20.0144mm.

1. Find a 50% confidence interval for the true mean μ.
2. Find an 80% confidence interval for the true mean μ.
3. Find a 90% confidence interval for the true mean μ.
4. Find a 95% confidence interval for the true mean μ.
5. Find a 99% confidence interval for the true mean μ.
6. Based on these intervals, does it seem likely that the true mean is actually 20.000mm?
7. If we wanted a 99% margin of error of 0.005mm for our estimate, how many bolts should be sampled?

Examples, V

Example: The diameters of bolts manufactured at Factory X are distributed normally with standard deviation 0.01mm. A random sample of 10 bolts from one lot has average diameter 20.0144mm.

$1-\alpha$	50%	80%	90%	95%	98%	99%	99.5%	99.9%
c	0.6745	1.2816	1.6449	1.9600	2.3263	2.5758	2.8070	3.2905

1. Find a 50% confidence interval for the true mean μ.

Examples, V

Example: The diameters of bolts manufactured at Factory X are distributed normally with standard deviation 0.01mm. A random sample of 10 bolts from one lot has average diameter 20.0144mm.

$1-\alpha$	50%	80%	90%	95%	98%	99%	99.5%	99.9%
c	0.6745	1.2816	1.6449	1.9600	2.3263	2.5758	2.8070	3.2905

1. Find a 50% confidence interval for the true mean μ.

- Here, we have $n = 10$, $\hat{\mu} = 20.0144$mm, and $\sigma/\sqrt{n} = 0.01/\sqrt{10} = 0.00316$mm.
- Thus, the 50% confidence interval is $\hat{\mu} \pm 0.6745 \cdot \sigma/\sqrt{n} = (20.0123\text{mm}, 20.0165\text{mm})$.

2. Find an 80% confidence interval for the true mean μ.

Examples, V

Example: The diameters of bolts manufactured at Factory X are distributed normally with standard deviation 0.01mm. A random sample of 10 bolts from one lot has average diameter 20.0144mm.

$1-\alpha$	50%	80%	90%	95%	98%	99%	99.5%	99.9%
c	0.6745	1.2816	1.6449	1.9600	2.3263	2.5758	2.8070	3.2905

1. Find a 50% confidence interval for the true mean μ.

- Here, we have $n = 10$, $\hat{\mu} = 20.0144$mm, and $\sigma/\sqrt{n} = 0.01/\sqrt{10} = 0.00316$mm.
- Thus, the 50% confidence interval is $\hat{\mu} \pm 0.6745 \cdot \sigma/\sqrt{n} = (20.0123\text{mm}, 20.0165\text{mm})$.

2. Find an 80% confidence interval for the true mean μ.

- The 80% confidence interval is $\hat{\mu} \pm 1.2816 \cdot \sigma/\sqrt{n} = (20.0112\text{mm}, 20.0176\text{mm})$.

Examples, VI

Example: The diameters of bolts manufactured at Factory X are distributed normally with standard deviation 0.01mm. A random sample of 10 bolts from one lot has average diameter 20.0144mm.

$1-\alpha$	50%	80%	90%	95%	98%	99%	99.5%	99.9%
c	0.6745	1.2816	1.6449	1.9600	2.3263	2.5758	2.8070	3.2905

3. Find a 90% confidence interval for the true mean μ.

Examples, VI

Example: The diameters of bolts manufactured at Factory X are distributed normally with standard deviation 0.01mm. A random sample of 10 bolts from one lot has average diameter 20.0144mm.

$1-\alpha$	50%	80%	90%	95%	98%	99%	99.5%	99.9%
c	0.6745	1.2816	1.6449	1.9600	2.3263	2.5758	2.8070	3.2905

3. Find a 90% confidence interval for the true mean μ.

- $\hat{\mu} = 20.0144$mm and $\sigma/\sqrt{n} = 0.01/\sqrt{10} = 0.00316$mm.
- The 90% confidence interval is $\hat{\mu} \pm 1.6449 \cdot \sigma/\sqrt{n} = (20.0092\mathrm{mm}, 20.0196\mathrm{mm})$.

4. Find an 95% confidence interval for the true mean μ.

Examples, VI

Example: The diameters of bolts manufactured at Factory X are distributed normally with standard deviation 0.01mm. A random sample of 10 bolts from one lot has average diameter 20.0144mm.

$1-\alpha$	50%	80%	90%	95%	98%	99%	99.5%	99.9%
c	0.6745	1.2816	1.6449	1.9600	2.3263	2.5758	2.8070	3.2905

3. Find a 90% confidence interval for the true mean μ.
   - $\hat{\mu} = 20.0144\text{mm}$ and $\sigma/\sqrt{n} = 0.01/\sqrt{10} = 0.00316\text{mm}$.
   - The 90% confidence interval is $\hat{\mu} \pm 1.6449 \cdot \sigma/\sqrt{n} = (20.0092\mathrm{mm}, 20.0196\mathrm{mm})$.
4. Find an 95% confidence interval for the true mean μ.
   - The 95% confidence interval is $\hat{\mu} \pm 1.9600 \cdot \sigma/\sqrt{n} = (20.0082\mathrm{mm}, 20.0206\mathrm{mm})$.

Examples, VII

Example: The diameters of bolts manufactured at Factory X are distributed normally with standard deviation 0.01mm. A random sample of 10 bolts from one lot has average diameter 20.0144mm.

$1-\alpha$	50%	80%	90%	95%	98%	99%	99.5%	99.9%
c	0.6745	1.2816	1.6449	1.9600	2.3263	2.5758	2.8070	3.2905

5. Find a 99% confidence interval for the true mean μ.

Examples, VII

Example: The diameters of bolts manufactured at Factory X are distributed normally with standard deviation 0.01mm. A random sample of 10 bolts from one lot has average diameter 20.0144mm.

$1-\alpha$	50%	80%	90%	95%	98%	99%	99.5%	99.9%
c	0.6745	1.2816	1.6449	1.9600	2.3263	2.5758	2.8070	3.2905

5. Find a 99% confidence interval for the true mean μ.

- The 99% confidence interval is $\hat{\mu} \pm 2.5758 \cdot \sigma/\sqrt{n} = (20.0063\text{mm}, 20.0225\text{mm})$.

6. Based on these intervals, does it seem likely that the true mean is actually 20.000mm?

Examples, VII

Example: The diameters of bolts manufactured at Factory X are distributed normally with standard deviation 0.01mm. A random sample of 10 bolts from one lot has average diameter 20.0144mm.

$1-\alpha$	50%	80%	90%	95%	98%	99%	99.5%	99.9%
c	0.6745	1.2816	1.6449	1.9600	2.3263	2.5758	2.8070	3.2905

5. Find a 99% confidence interval for the true mean μ.

- The 99% confidence interval is $\hat{\mu} \pm 2.5758 \cdot \sigma/\sqrt{n} = (20.0063\text{mm}, 20.0225\text{mm})$.

6. Based on these intervals, does it seem likely that the true mean is actually 20.000mm?

- Even the 99% confidence interval does not contain 20.000mm.
- So it does not seem very likely that the true mean is actually 20.000mm, since we would only expect the population mean to be outside our 99% confidence interval 1% of the time.

Examples, VIII

A few comments about this last observation:

- Notice that that the average diameter of the bolts in the sample only differs from the desired one by 0.0144mm, which was 1.44 times the standard deviation of the bolt diameter.
- Nevertheless, based on our confidence intervals, this is in fact very strong evidence that the true mean of this lot of bolts is not actually 20mm.
- On Thursday, we will extend this type of analysis to describe methods for testing the hypothesis that the bolt diameter is actually equal to 20mm. (But it is worth thinking right now about how you might try to do this!)

Examples, IX

Example: The diameters of bolts manufactured at Factory X are distributed normally with standard deviation 0.01mm. A random sample of 10 bolts from one lot has average diameter 20.0144mm.

$1-\alpha$	50%	80%	90%	95%	98%	99%	99.5%	99.9%
c	0.6745	1.2816	1.6449	1.9600	2.3263	2.5758	2.8070	3.2905

7. If we wanted a 99% margin of error of 0.005mm for our estimate, how many bolts should be sampled?

Examples, IX

Example: The diameters of bolts manufactured at Factory X are distributed normally with standard deviation 0.01mm. A random sample of 10 bolts from one lot has average diameter 20.0144mm.

$1-\alpha$	50%	80%	90%	95%	98%	99%	99.5%	99.9%
c	0.6745	1.2816	1.6449	1.9600	2.3263	2.5758	2.8070	3.2905

7. If we wanted a 99% margin of error of 0.005mm for our estimate, how many bolts should be sampled?

- The margin of error for a 99% confidence interval in this setting is $2.5758 \cdot \sigma/\sqrt{n}$.
- Since this quantity is required to be 0.005mm, solving for n gives $n = \left(\frac{2.5758 \cdot 0.01\text{mm}}{0.005\text{mm}}\right)^2 \approx 26.54$.
- This means a sample of 27 bolts would be sufficient to give the desired precision.

Examples, X

Example: A marine biologist measures the lengths of 100 adult blue whales. In his sample, the average length was 27.11m with a standard deviation of 1.3m. Assuming that this standard deviation is correct for the full population, find:

1. A 98% confidence interval for the average length of a blue whale.
2. The number of blue whales that would need to be measured to give a 98% confidence interval with half the margin of error as the one just found.
3. The probability that if another 100 blue whales were independently sampled, both 98% confidence intervals would contain the true mean.

Examples, XI

<u>Example</u>: A marine biologist measures the lengths of 100 adult blue whales. In his sample, the average length was 27.11m with a standard deviation of 1.3m. Assuming that this standard deviation is correct for the full population, find:

$1-\alpha$	50%	80%	90%	95%	98%	99%	99.5%	99.9%
c	0.6745	1.2816	1.6449	1.9600	2.3263	2.5758	2.8070	3.2905

1. A 98% confidence interval for the average length of a blue whale.

Examples, XI

Example: A marine biologist measures the lengths of 100 adult blue whales. In his sample, the average length was 27.11m with a standard deviation of 1.3m. Assuming that this standard deviation is correct for the full population, find:

$1-\alpha$	50%	80%	90%	95%	98%	99%	99.5%	99.9%
c	0.6745	1.2816	1.6449	1.9600	2.3263	2.5758	2.8070	3.2905

1. A 98% confidence interval for the average length of a blue whale.

- For (i), we have $n = 100$, $\hat{\mu} = 27.11\text{m}$, and $\sigma/\sqrt{n} = 1.3\text{m}/\sqrt{100} = 0.13\text{m}$.
- Thus, using the table, we obtain the 98% confidence interval $(\hat{\mu} - 2.3263 \cdot \sigma/\sqrt{n},\ \hat{\mu} + 2.3263 \cdot \sigma/\sqrt{n}) = (26.81\text{m}, 27.41\text{m})$.

Examples, XII

<u>Example</u>: A marine biologist measures the lengths of 100 adult blue whales. In his sample, the average length was 27.11m with a standard deviation of 1.3m. Assuming that this standard deviation is correct for the full population, find:

$1-\alpha$	50%	80%	90%	95%	98%	99%	99.5%	99.9%
c	0.6745	1.2816	1.6449	1.9600	2.3263	2.5758	2.8070	3.2905

2. The number of blue whales that would need to be measured to give a 98% confidence interval with half the margin of error as the one just found.

Examples, XII

Example: A marine biologist measures the lengths of 100 adult blue whales. In his sample, the average length was 27.11m with a standard deviation of 1.3m. Assuming that this standard deviation is correct for the full population, find:

$1-\alpha$	50%	80%	90%	95%	98%	99%	99.5%	99.9%
c	0.6745	1.2816	1.6449	1.9600	2.3263	2.5758	2.8070	3.2905

2. The number of blue whales that would need to be measured to give a 98% confidence interval with half the margin of error as the one just found.

- The width of the confidence interval is $2.3263 \cdot \sigma/\sqrt{n}$.
- Thus, to halve the width we would require a value of n four times as large, which is $n = 400$.

Examples, XIII

Example: A marine biologist measures the lengths of 100 adult blue whales. In his sample, the average length was 27.11m with a standard deviation of 1.3m. Assuming that this standard deviation is correct for the full population, find:

3. The probability that if another 100 blue whales were independently sampled, both 98% confidence intervals would contain the true mean.

Examples, XIII

Example: A marine biologist measures the lengths of 100 adult blue whales. In his sample, the average length was 27.11m with a standard deviation of 1.3m. Assuming that this standard deviation is correct for the full population, find:

3. The probability that if another 100 blue whales were independently sampled, both 98% confidence intervals would contain the true mean.

- By definition, each confidence interval has a probability 0.98 of containing the true mean.
- Since these samples are independent, the probability that both contain the true mean is $(0.98)^2 = 0.9604$.

Examples, XIV

We will remark that the assumption in the previous example, that the sample standard deviation is equal to the population standard deviation, is not generally valid in practice. (I included it primarily so I could make this very point right now!)

- Indeed, as we have already seen, the sample variance is not an unbiased estimator of the population variance.
- If we employ Bessel's correction (i.e., compute the sample standard deviation with denominator $n - 1$ rather than n) then we will obtain a better estimate of the population standard deviation.

Examples, XV

- More precisely, our discussion effectively analyzes the ratio $\dfrac{x - \overline{x}}{\sigma/\sqrt{n}}$ where σ is the (known) population standard deviation by observing that this random variable has a standard normal distribution.
- However, if we replace the population standard deviation σ by the sample standard deviation S, the resulting ratio $\dfrac{x - \overline{x}}{S/\sqrt{n}}$ is no longer normally distributed.
- We can therefore not construct confidence intervals using the procedure described above.
- As we will discuss next week, the random variable $\dfrac{x - \overline{x}}{S/\sqrt{n}}$ actually follows a distribution known as the t distribution.

Examples, XVI

Example: The weight of domestic housecats is normally distributed with a standard deviation of 0.24kg. Some cats from two feral colonies, Colony A and Colony B, are each weighed. The 16 cats from Colony A had an average weight of 3.95kg while the 9 cats from Colony B had an average weight of 4.24kg.

1. Find a 90% CI for the average weight of cats from Colony A.
2. Find a 90% CI for the total weight of 16 cats from Colony A.
3. Find a 90% CI for the average weight of cats from Colony B.
4. Find a 90% CI for the total weight of 9 cats from Colony B.
5. Find a 90% CI for the difference of the average weights of the two colonies.

$1-\alpha$	50%	80%	90%	95%	98%	99%	99.5%	99.9%
c	0.6745	1.2816	1.6449	1.9600	2.3263	2.5758	2.8070	3.2905

Examples, XVII

Example: The weight of domestic housecats is normally distributed with a standard deviation of 0.24kg. The 16 cats from Colony A had an average weight of 3.95kg.

1. Find a 90% CI for the average weight of cats from Colony A.

Examples, XVII

Example: The weight of domestic housecats is normally distributed with a standard deviation of 0.24kg. The 16 cats from Colony A had an average weight of 3.95kg.

1. Find a 90% CI for the average weight of cats from Colony A.

- For Colony A, we have $n = 16$, $\hat{\mu} = 3.95\text{kg}$, and $\sigma/\sqrt{n} = 0.06\text{kg}$. For $\alpha = 0.90$ we have $c = 1.6449$.
- Thus, we have a 90% confidence interval for the average weight given by $\hat{\mu} \pm 1.6449\sigma/\sqrt{n} = (3.85\text{kg}, 4.05\text{kg})$.

Examples, XVIII

<u>Example</u>: The weight of domestic housecats is normally distributed with a standard deviation of 0.24kg. The 16 cats from Colony A had an average weight of 3.95kg.

2. Find a 90% CI for the total weight of 16 cats from Colony A.

Examples, XVIII

Example: The weight of domestic housecats is normally distributed with a standard deviation of 0.24kg. The 16 cats from Colony A had an average weight of 3.95kg.

2. Find a 90% CI for the total weight of 16 cats from Colony A.

- Since there are 16 cats, we simply scale the interval we just found by 16.
- This yields the 90% confidence interval $(61.6\mathrm{kg}, 64.8\mathrm{kg})$.
- Alternatively, we could observe that the total weight is normally distributed with mean $16 \cdot 3.95 = 63.2$kg and standard deviation $\sqrt{16} \cdot 0.24 = 0.96$kg.
- Then the desired confidence interval is $63.2 \pm 1.6449 \cdot 0.96$kg, which is the same as above.

Examples, XIX

Example: The weight of domestic housecats is normally distributed with a standard deviation of 0.24kg. The 9 cats from Colony B had an average weight of 4.24kg.

3. Find a 90% CI for the average weight of cats from Colony B.

Examples, XIX

<u>Example</u>: The weight of domestic housecats is normally distributed with a standard deviation of 0.24kg. The 9 cats from Colony B had an average weight of 4.24kg.

3. Find a 90% CI for the average weight of cats from Colony B.

- Here, we have $n = 9$, $\hat{\mu} = 4.24\mathrm{kg}$, and $\sigma/\sqrt{n} = 0.08\mathrm{kg}$.
- Thus, we have a 90% confidence interval for the average weight given by $\hat{p} \pm 1.6449\sigma/\sqrt{n} = (4.11\mathrm{kg}, 4.37\mathrm{kg})$.

4. Find a 90% CI for the total weight of 9 cats from Colony B.

Examples, XIX

<u>Example</u>: The weight of domestic housecats is normally distributed with a standard deviation of 0.24kg. The 9 cats from Colony B had an average weight of 4.24kg.

3. Find a 90% CI for the average weight of cats from Colony B.

- Here, we have $n = 9$, $\hat{\mu} = 4.24\text{kg}$, and $\sigma/\sqrt{n} = 0.08\text{kg}$.
- Thus, we have a 90% confidence interval for the average weight given by $\hat{p} \pm 1.6449\sigma/\sqrt{n} = (4.11\text{kg}, 4.37\text{kg})$.

4. Find a 90% CI for the total weight of 9 cats from Colony B.

- Since there are 9 cats, we simply scale the interval we just found by 9, yielding $(37.0\text{kg}, 39.3\text{kg})$.

Examples, XX

Example: The weight of domestic housecats is normally distributed with a standard deviation of 0.24kg. The 16 cats from Colony A had an average weight of 3.95kg, while the 9 cats from Colony B had an average weight of 4.24kg.

5. Find a 90% CI for the difference of the average weights of the two colonies.

Examples, XX

Example: The weight of domestic housecats is normally distributed with a standard deviation of 0.24kg. The 16 cats from Colony A had an average weight of 3.95kg, while the 9 cats from Colony B had an average weight of 4.24kg.

5. Find a 90% CI for the difference of the average weights of the two colonies.

- This one is a bit trickier. The idea is to use the fact that the difference between two independent, normally distributed random variables is also normally distributed.
- Specifically, suppose A is normally distributed with mean μ_A and standard deviation σ_A, while B is normally distributed with mean μ_B and standard deviation σ_B.
- Then $B - A$ is normally distributed with mean $\mu_B - \mu_A$ and standard deviation $\sqrt{\sigma_A^2 + \sigma_B^2}$.

Examples, XXI

Example: The weight of domestic housecats is normally distributed with a standard deviation of 0.24kg. The 16 cats from Colony A had an average weight of 3.95kg, while the 9 cats from Colony B had an average weight of 4.24kg.

5. Find a 90% CI for the difference of the average weights of the two colonies.

- That means the difference in the average weights is normally distributed with mean $4.24\text{kg} - 3.95\text{kg} = 0.29\text{kg}$ and standard deviation $\sqrt{(0.06\text{kg})^2 + (0.08\text{kg})^2} = 0.1\text{kg}$.
- Thus, by our discussion, the 90% confidence interval for the difference in the average weights will be $0.29\text{kg} \pm 1.6449 \cdot 0.1\text{kg} = (0.125\text{kg}, 0.454\text{kg})$.

What Do Confidence Intervals Feel Like?, I

We have given a formal definition of confidence intervals, which works well enough. However, it is also important to get a bit of an intuitive sense of what confidence intervals actually feel like.

So let's do a brief experiment, as follows:

What Do Confidence Intervals Feel Like?, I

We have given a formal definition of confidence intervals, which works well enough. However, it is also important to get a bit of an intuitive sense of what confidence intervals actually feel like.

So let's do a brief experiment, as follows:

- I will give 10 questions, each of which has a precise numerical answer.
- For each item, you will write down a 90% confidence interval.
- Remember: this is an interval around your estimate that you expect to contain the true value 90% of the time.
- I will then give you the actual values and you will tally how many values actually landed in your confidence interval.
- The point is not to look up any of the values (this isn't a trivia contest!): use only your vague sense of what the answers might be, and then quantify your uncertainty.

What Do Confidence Intervals Feel Like?, II

Questions go here.

What Do Confidence Intervals Feel Like?, IV

All right... now I'll give you the answers.

What Do Confidence Intervals Feel Like?, IV

All right... now I'll give you the answers.

Answers go here.

What Do Confidence Intervals Feel Like?, V

It usually happens that most people vastly overestimate their ability to construct intervals with 90% confidence[1].

- Remember that the idea of a 90% confidence interval is that, if you construct many of them, about 90% of them should contain the actual parameter value.
- If you did get exactly 9 out of 10, great!
- But most people find that they are actually giving more like a 50% confidence interval here.

[1] Insert obvious joke regarding overconfidence about confidence here.

What Do Confidence Intervals Feel Like?, VI

Let's try another round. This time, give both a 50% confidence interval and a 90% confidence interval for each item.

What Do Confidence Intervals Feel Like?, VII

Questions go here.

What Do Confidence Intervals Feel Like?, VII

All right... now I'll give you the answers.

What Do Confidence Intervals Feel Like?, VII

All right... now I'll give you the answers.

Answers go here.

What Do Confidence Intervals Feel Like?, VIII

So, aside from putting together some fun trivia questions, the purpose of this exercise is to help you understand what a confidence interval is really measuring, and what the percentage actually represents.

Specifically: if everyone in the class writes down a 90% confidence interval for an unknown quantity like the ones we just went through, then overall we should expect about 90% of the confidence intervals to contain the true value.

Summary

We introduced interval estimation and defined confidence intervals.

We discussed how to construct confidence intervals for normally-distributed variables, and gave several examples.

We discussed how to interpret confidence intervals.

Next lecture: Interval estimation and confidence intervals (part 2)