Introduction to
Uncertainty in Physical Measurements
by David Harrison and Jason Harlow

Last updated Sep. 24, 2015 by Jason Harlow.
This is loosely based on a wonderful series of documents written by David M. Harrison, Department of Physics, University of Toronto in 2014. These documents are available at: http://www.upscale.utoronto.ca/PVB/Harrison/GUM/

1. Introduction

Almost every time you make a measurement, the result will not be an exact number, but it will be a *range* of possible values. The range of values associated with a measurement is described by the **uncertainty**. The uncertainty is a number which follows the ± sign. For example, in the measurement (8 ± 2), 8 is the value, and 2 is the uncertainty. Since all of science depends on measurements, it is important to understand uncertainties and get used to using them.

[An exception to this rule is when you are counting small numbers, such as a pile of 3 apples. "Three" is an exact number of apples that does not have an uncertainty. However, if you estimate the number of apples in the back of a large apple truck and find there are 1600, this number does have an uncertainty!]

The uncertainty in a measurement is sometimes called the "error". This is an outdated term, because the word "error" implies that some kind of a mistake has been made. On the contrary, uncertainty is a necessary part of any measurement, and it would be a mistake not to report it!

We should emphasize right now that a "correct" experiment is one that has been correctly performed. *You do **not** determine the uncertainty in an experimentally measured quantity by comparing it to some number found in a book or web page!* You determine each uncertainty as part of your experiment; the value and the uncertainty are two numbers that are measured together, and you report them both, independently of what anyone else may have reported in the past.

Also, although we will be exploring mathematical and statistical procedures that are used to determine the uncertainty in an experimentally measured quantity, as you will see these are often just "rules of thumb" and sometimes a good experimentalist uses his or her intuition and common sense to simply guess what the uncertainty is. That is okay sometimes!

Example 1: Drawing a Histogram of repeated measurements
Imagine you have a cart on a track with a fan attached to it which causes it to accelerate along the track. You release the cart from rest and then use a digital stopwatch to measure the time it takes the cart to travel 1.50 m. You estimate the uncertainty in the distance you measured to be 1 cm, so actually the distance is 1.50 ± 0.01 m. You then repeat the time measurements for a total of 30 trials. The time measurements are shown in the Table 1.

Trial #	Time (s)
1	5.55
2	5.50
3	5.56
4	5.51
5	5.49
6	5.46
7	5.49
8	5.51
9	5.32
10	5.50

Trial #	Time (s)
11	5.49
12	5.45
13	5.30
14	5.67
15	5.64
16	5.61
17	5.66
18	5.69
19	5.38
20	5.37

Trial #	Time (s)
21	5.74
22	5.54
23	5.56
24	5.47
25	5.49
26	5.54
27	5.45
28	5.59
29	5.52
30	5.45

Table 1. Time for a cart to start from rest, accelerate at a constant rate, and travel a distance of 1.50 ± 0.01 m. The measurement was repeated 30 times.

a. Make a histogram of the results. On the horizontal axis, plot the time from 5.0 to 6.0 seconds with tick-marks at intervals of 0.05 seconds. On the vertical axis, plot the number of trials that fall into each range.
b. What is the height of the histogram? [number of trials in the most populated bin]
c. What is the width of the histogram? [size of the central part of the range which contains about 2/3 of the trials.]
d. Sketch a smooth bell-shaped curve through your histogram to approximate the tops of the histogram. What is the centre of this curve? [this approximates the mean].

Example 1: Answers

a. See the Excel spreadsheet, first sheet, available for download along with this document for work. Here is a plot of the histogram:

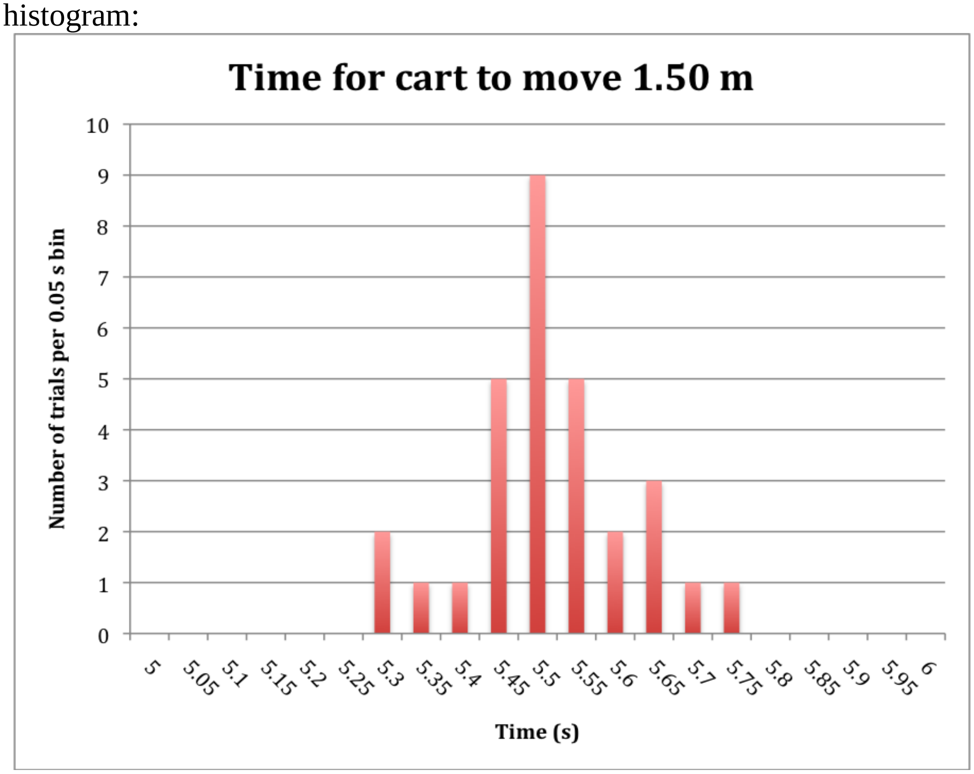

b. The height of the histogram is 9.
c. 2/3 of the 30 trials is 20. In the three bins with centres at 5.45, 5.5 and 5.55, there are a total of 5 + 9 + 5 = 19 trials, which is almost 2/3 of the total. Therefore, I would say the width of the histgram is about 3 bins, or 5.55 – 5.45 = 0.1 seconds.
d. The smooth curve was not done in Excel, but if I was to do it by hand on a print-out, it would probably have its centre slightly to the right of 5.5 s, because the tail of this histogram extends further toward the right than the left. So perhaps the mean is about 5.51 s.

2. Normal Distribution

Each time you make a measurement, there is some probability that you will get a certain answer. A probability distribution is a curve which describes what the probability is for various measurements. The most important and widely used probability distribution is called the **Normal Distribution**. It was first popularized by the German mathematician Carl Friedrich Gauss in the early 1800s. It is also sometimes called the Gaussian distribution, or the bell-curve. The formula is:

$$N(x) = Ae^{\frac{-(x-\bar{x})^2}{2\sigma^2}}$$

which looks like:

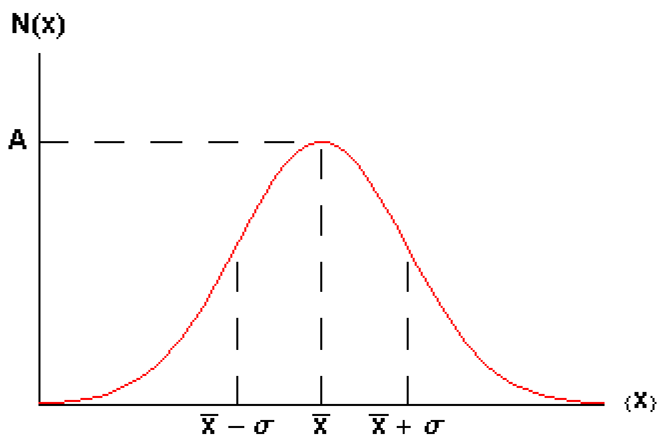

The symbol A is called the maximum amplitude.

The symbol $\bar{x}$ is called the **mean** or average.

The symbol σ is called the **standard deviation** of the distribution. Statisticians often call the square of the standard deviation, σ^2, the variance; we will not use that name. Note that σ is a measure of the width of the curve: a larger σ means a wider curve. (σ is the lower case Greek letter sigma.)

Note that 68% of the area under the curve of a Gaussian lies between the mean minus the standard deviation and the mean plus the standard deviation. Similarly, 95% of the curve is between the mean minus twice the standard deviation and the mean plus twice the standard deviation.

3. Using the Gaussian

Suppose you make *N* measurements of a quantity *x*, and you expect these measurements to be normally distributed. Each measurement, or trial, you label with a number *i*, where *i* = 1, 2, 3, etc. You do not know what the true mean of the distribution is, and you cannot know this. However, you can estimate the mean by adding up all the individual measurements and dividing by *N*:

Estimate of Mean: $$\bar{x}_{\text{est}} = \frac{1}{N}\sum_{i=1}^{N} x_i$$

Similarly, it is impossible to know the true standard deviation of the distribution. However, we can estimate the standard deviation using our *N* measurements. The best estimate of the standard deviation is:

Estimate of Standard Deviation: $$\sigma_{\text{est}} = \sqrt{\frac{1}{N-1}\sum_{i=1}^{N}(x_i - \bar{x}_{\text{est}})^2}$$

The quantity $N-1$ is called the *number of degrees of freedom*. In the case of the standard deviation estimate, it is the number of measurements minus one because you used one number from a previous calculation (mean) in order to find the standard deviation.

Example 2: Consider the 30 measurements of time in the Table 1.

a. What is the estimated mean?
b. What is the estimated standard deviation?
c. How many of these 30 measurements fall in the range $\bar{x}_{\text{est}} \pm \sigma_{\text{est}}$? What percentage is that?

Example 2: Answers

a. See the Excel spreadsheet, second sheet, available for download along with this document for work. The estimated mean is found to be 5.517 s.
b. The estimated standard deviation is found to be 0.103 s.
c. 21 of the 30 measurements fall within the range 5.517 ± 0.103 s, or $5.414 < t < 5.620$ s. This is 70% of the trials. That is pretty close to 68%, as expected.

There is roughly a 68% chance that any measurement of a sample taken at random will be within one standard deviation of the mean. Usually the mean is what we wish to know and each individual measurement almost certainly differs from the true value of the mean by some amount. But there is a 68% chance that any single measurement lies with one standard deviation of this true value of the mean. Thus it is reasonable to say that:

The standard deviation is the uncertainty in each individual measurement of the sample.

The uncertainty in a quantity is usually indicated by a *u*, so the above statement may be written as $u_{xi} = \sigma$. This uncertainty is often called *statistical*. We shall see another type of uncertainty in the next section.

4. Reading Uncertainty

In the previous section we saw that when we repeat measurements of some quantity in which random statistical factors lead to a spread in the values from one trial to the next, it is reasonable to set the uncertainty in each individual measurement equal to the standard deviation of the sample. Here we discuss another uncertainty that arises when we do a direct measurement of some quantity: the reading uncertainty.

For example, imagine you use a metric ruler to measure the length of a crayon. You line up the bottom end of the crayon with 0 cm, and the image below shows what you see for the tip of the crayon. The crayon appears to be about 9.32 cm long. What is the reading uncertainty?

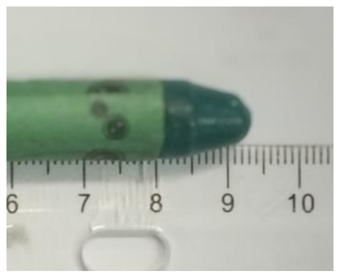

To determine the reading uncertainty in this measurement we have to answer the question: what is the minimum and maximum values that the position could have for which we will not see any difference? There is no fixed rule that will allow us to answer this question. Instead we must use our *intuition* and *common sense*!

Could the crayon actually be as long as 9.4 cm? No, I don't think so. It's clearly shorter than that. Could it be 9.35 cm? Maybe. And it could be as short as 9.30 cm, but, in my opinion, no shorter. So the range is about 9.30 to 9.35 cm. A reasonable estimate of the reading uncertainty of this measurement is half the range: ± 0.025 cm, which, to be cautious, we might round up to 0.03 cm. Then we would state that the length of the crayon is 9.32 ± 0.03 cm.

For your eyes you may wish to instead associate a reading uncertainty of 0.02 cm with the position; this is also a reasonable number. A reading uncertainty of 0.04 cm, though, is probably too pessimistic. And a reading uncertainty much less than 0.01 cm is probably too optimistic.

We assume that the reading uncertainty indicates a spread in repeated measurements, just like the standard deviation discussed in the previous section. So if we get a collection of objective observers together to look at the crayon above, we expect most (ie more than 68%) of all observers will report a value between 9.29 and 9.35 cm.

Note that there is often a trade-off when assigning a reading uncertainty such as above. On the one hand we want the uncertainty to be as small as possible, indicating a precise measurement. However we also want to insure that measured value probably lies within uncertainties of the "true" value of the quantity, which means we don't want the uncertainty to be too small.

For a measurement with an instrument with a digital readout, the reading uncertainty is "± one-half of the last digit."

We illustrate with a digital thermometer shown to the right. The phrase "± one-half of the last digit" above is the language commonly used in manufacturer's specification sheets of their instruments. It can be slightly misleading. It does not mean one half of the value of the last digit, which for this thermometer photo is 0.4°C. It means one-half of the power of ten represented in the last digit. Here, the last digit represents values of a tenth of a degree, so the reading uncertainty is 1/2 x 0.1 = 0.05°C. You should write the temperature as 12.80 ± 0.05 °C.

It may be confusing to notice that we have two different specifications for the uncertainty in a directly measured quantity: the **standard deviation** and the **reading uncertainty**. Both are indicators of a spread in the values of repeated measurements. They both are describing the precision of the measurement. So you may ask: *What is the uncertainty in the quantity?*

The answer is both. But fortunately it is almost always the case that one of the two is much larger than the other, and in this case we choose the larger to be the uncertainty.

For example, if every time you measure something you always get the same numerical answer, this indicates that the reading uncertainty is dominant.

However, if every time you measure something you get different answers which differ more than the reading uncertainty you might estimate, then the standard deviation is dominant. In these cases you can usually ignore the reading uncertainty. For example, consider the rolling fan-cart we looked at in Table 1. The reading uncertainty of the digital stopwatch is ± 0.005 seconds, which is negligible compared to the standard deviation, which we found to be 0.1 seconds. Thus here the uncertainty in each measurement is the standard deviation.

Often just thinking for a moment in advance about a measurement will tell you whether repeating it will be likely to show a spread of values not accounted for by the reading uncertainty. This in turn will tell you whether you need to bother repeating it at all. If you don't know in advance whether or not you need to repeat a measurement you can usually find out by repeating it three or four times.

5. Significant Figures

Most introductory Physics textbooks (including the one chosen for this course) include a little section on significant figures. It provides some rules for how many significant figures to include if your number does not include an uncertainty. However, in experimental science *every number you measure or compute should include an uncertainty!* Therefore, the rules in your textbook are useless for any actual numbers measured in the Practicals. In this section, we will teach you the rules for significant figures when there is an uncertainty associated with every value.

Let's consider again the 30 data points in Table 1. The estimated standard deviation is numerically equal to 0.102933971 seconds, which is larger than the reading uncertainty for these measurements. (By "numerically" we mean that is what the calculator read when we computed the standard deviation, or what Excel produced when we used the STDEV function.) Since the estimated standard deviation is larger than the reading uncertainty, it will be the uncertainty in the value of each of the data points.

Consider one of these data points, say, the 5^{th} trial, for which we measured 5.49 seconds. If you take the estimated standard deviation to be 0.102933971, then the data point has a value of 5.49 ± 0.102933971. What this means is that there is about a 68% chance that the true value is somewhere between 5.387066029 and 5.592933971 seconds.

Keep in mind that all of these are estimates, as we have only 30 data points, not infinity. Also, 68% is a somewhat arbitrary probability, which we have chosen to represent "most of the time". Clearly we are using WAY too many significant figures here! It would be just as instructive to say that there is about a 68% chance that the true value is somewhere between 5.4 and 5.6 seconds. Or, you could say the measurement is 5.5 ± 0.1 s. In fact it is not only more concise to report this, but it is more honest.

This example illustrates two general rules for significant figures used in experimental sciences:

1. **Uncertainties should be specified to one or two significant figures.**
2. **The most precise column in the number for the uncertainty should also be the most precise column in the number for the value.**

So if the uncertainty is specified to the 1/100th column, the quantity itself should also be specified to the 1/100th column.

Example 3. Express the following quantities to the correct number of significant figures:

a. 25.052 ± 1.502
b. 92 ± 3.14159
c. 0.0530854 ± 0.012194

Example 3: Answers

a. 25.1 ± 1.5 (or 25 ± 2 is also acceptable)
b. 92 ± 3
c. 0.053 ± 0.012 (or 0.05 ± 0.01 is also acceptable)

6. Propagation of Uncertainties of Precision

When you have two or more quantities with known uncertainties you may sometimes want to combine them to compute a derived number. If your uncertainties come from standard deviation and are not reading uncertainties, it is usually best to compute the derived number several times and use compute its standard deviation directly. However, if some reading uncertainties are larger than the standard deviation, you can use the rules of Uncertainty Propagation to infer the uncertainty in the derived quantity.

We assume that the two directly measured quantities are x and y, with uncertainties u_x and u_y respectively. The measurements x and y must be independent of each other.

The fractional uncertainty is the value of the uncertainty divided by the value of the quantity: u_x / x. The fractional uncertainty multiplied by 100 is the percentage uncertainty. Everything is this section assumes that the uncertainty is "small" compared to the value itself, i.e. that the fractional uncertainty is much less than one.

For many situations, we can find the uncertainty in the result z using three simple rules:

Rule # 1 (sum or difference rule):

If $z = x + y$

Or $z = x - y$

Then $$u_z = \sqrt{u_x^2 + u_y^2}$$

Rule #2 (product or division rule):

If $z = xy$

Or $z = x/y$

Then $$\frac{u_z}{z} = \sqrt{\left(\frac{u_x}{x}\right)^2 + \left(\frac{u_y}{y}\right)^2}$$

Note that by this rule, if x is an exact number ($u_x = 0$), then $z = xy$ means $u_z = |x|\, u_y$. The uncertainty is simply multiplied by the magnitude of the same exact number.

Rule #3 (exponent rule):

If $z = x^n$

Then $$\frac{u_z}{z} = n\frac{u_x}{x}$$

Example 4. As with Example 1, you have a cart on a track with a fan attached to it, which causes it to accelerate along the track. You release the cart from rest and then use a digital stopwatch to measure the time it takes the cart to travel 1.50 ± 0.01 m. You then repeat the time measurements for a total of 30 trials. The time measurements are shown in the Table 1. You model the distance as a function of time with the kinematic equation $d = ½\, a\, t^2$. From your measurements of d and t you wish to derive a, which is $a = 2d/t^2$.

a. Compute individual estimates of acceleration for each time measurement, using d = 1.50 exactly. Estimate the mean and standard deviation of all these values of a.

b. As discussed in Section 5, the best way to report the measurement for the 5th trial is: $t = 5.5 \pm 0.1$ s. Combine this with $d = 1.50 \pm 0.01$ m and propagate uncertainties to compute acceleration, *a*.

Example 4: Answers

a. The mean of the acceleration values is 0.0987 m/s^2, and the standard deviation is 0.0037 m/s^2.
b. The equation is $a = 2d/t^2$, where 2 is an exact number (with no uncertainty). To compute the uncertainty, first use Rule#3 to find the uncertainty in $z = t^2$: $u_z = 2z(u_t/t) = 2(5.5^2)(0.1/5.5) = 1.1$. Then use Rule#2 to get the uncertainty in $y = d/t^2 = d/z$:
$u_y = y\sqrt{\left(\frac{u_d}{d}\right)^2 + \left(\frac{u_z}{z}\right)^2} = 0.0496\sqrt{\left(\frac{0.01}{1.5}\right)^2 + \left(\frac{1.1}{30.25}\right)^2}$ = 0.0018. Then use Rule#2 again with the exact number 2 as a multiplier, which simply increases the uncertainty by 2: $u_a = 0.0037$ m/s^2. It should not be surprising to you that the uncertainty in one measurement of acceleration should be equal to the standard deviation we computed in part a.

7. The Uncertainty in the Mean

We have seen that when the data have uncertainties of precision we may only estimate the value of the mean. We are now ready to find the uncertainty in this estimate of the mean.

Recall that to calculate the estimated mean we use:

$$\bar{x}_{\text{est}} = \frac{1}{N}\sum_{i=1}^{N} x_i$$

Each individual measurement x_i has the same uncertainty, u_x, which is usually the estimated standard deviation. To calculate the uncertainty in the sum of all the x_i, we use Uncertainty Propagation Rule #1 to find the uncertainty in the sum as $\sqrt{u_x^2 + u_x^2 + \ldots + u_x^2} = \sqrt{N u_x^2} = \sqrt{N}u_x$. We then use Rule #2 to find the uncertainty in the sum divided by the exact number *N*, which gives:

$$u_{\bar{x}est} = \frac{u_x}{\sqrt{N}}$$

Example 5.

a. What is the mean and the uncertainty in the mean for the 30 time measurements in Table 1? [Note that the uncertainty in the mean is ***not*** the same as the standard deviation!]
b. What is the mean and the uncertainty in the mean for the acceleration of the cart?

Example 5: Answers

a. 5.517 ± 0.019 s.
b. 9.867 ± 0.067 cm/s^2.

8. Practice Problems

1. A correctly performed experiment measures the flow rate of water from a tap, and finds that is it 90 ± 12 millilitres per second. A second correctly performed experiment uses a different method and finds that the flow rate is 110 ± 1 millilitres per second. What is the time, in s, for water from the tap to fill a 1.00 litre container? (There are 1000 millitres in a litre.)

2. A laboratory technician tells you the results of your chest X-ray are that you have a spherical tumor in your left lung that has a diameter of 3.4 mm, and the uncertainty of this measurement is ±0.4 mm. This uncertainty is the 1-sigma uncertainty based on the assumption of a normal distribution. What does this mean?

A. The tumor has a diameter between 3.0 mm and 3.8 mm.
B. There is about a 68% chance that the true diameter of the tumor lies between 3.0 mm and 3.8 mm.
C. There is about a 95% chance that the true diameter of the tumor lies between 3.0 mm and 3.8 mm.
D. The tumor has a diameter of either 3.0 mm or 3.8 mm.
E. The tumor has a diameter of 3.4 mm, but the technician may have made a mistake which makes it higher or lower by 0.4 mm.

3. Shown is a histogram of the body mass, in kg, of 3,000 sheep.

i. Roughly estimate the mean body mass.
ii. Roughly estimate the standard deviation of the distribution.

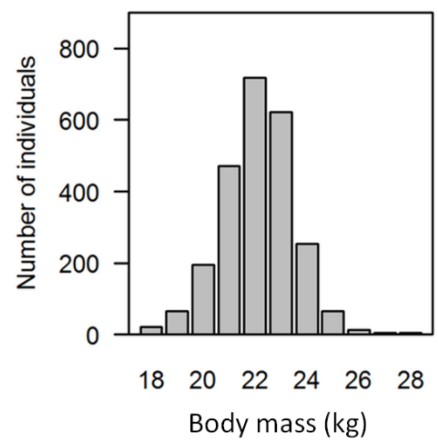

4. A satellite uses an infrared radiometer to indirectly measure the surface temperature of the ocean off the coast of British Columbia. The temperature is measured 5 times in the same week, and the measurements reported are 10.40, 10.35, 10.31, 10.56, and 10.43, all in degrees Celsius (°C). The estimated mean of these five measurements is 10.41°C. What is the uncertainty in this estimated mean?

5. Here is a table of the results of four repeated measurements of x in cm.

x_1	12.331
x_2	12.753
x_3	12.452
x_4	12.766

The mean of the four measurements is numerically 12.5755 cm (by numerically I mean the number as read on my calculator). The standard deviation of the four measurements is numerically 0.218196 cm. The first value x_1 has an uncertainty u_{x1}. What are the values of x_1 and u_{x1}? Be sure to use an appropriate number of significant figures in your answer.

6. An extremely accurate digital thermometer gives a reading of 37.2ºC when placed in a tissue sample. Repeated measurements give the same value of 37.2ºC. What is the uncertainty in this value?

7. You measure the distance between two trees to be 54.0 ± 0.5 m. Your friend, Bob, runs as fast as he can from the first tree to the second tree, touches it, and then runs all the way back to the first tree. You measure the time for his trip there and back to be 17.0 ± 0.2 s. What is Bob's average running speed over these 17 seconds?

8. An astronomer determines that the mass of the star Alpha Ceristinnae is 3.1 ± 0.1 times the mass of the sun. The mass of the sun is well-known, and is equal to $(1.98892 \pm 0.00025) \times 10^{30}$ kg. What is the mass of Alpha Ceristinnae, in kg?

9. You wish to measure the average speed of cars driving down St. George St. Over a distance of 125 ± 2 m you measure an average travel time of 11.3 ± 1.2 seconds. From these measurements you conclude that the average speed of cars is 11.1 m/s. What is the uncertainty in this value?

10. You wish to estimate the volume of one of the Practicals rooms. You model the volume of the room as $l \times w \times h$, and estimate that the length of the room is $l = 20 \pm 2$ m, the width is $w = 7 \pm 1$ m, and the distance from the floor to the ceiling is $h = 3.5 \pm 0.5$ m. From these estimates, what is that the total volume of the room?

11. A mass is suspended from a spring, and is free to oscillate up and down. Under the mass is a Motion Sensor which measures the *y*-position of the mass versus time. The plot shows the measurements, taken with a sampling rate of 100 measurements per second for 0.9 seconds. The amplitude, *A*, is defined as half the distance between the maximum height and minimum height of the mass. From these data, make a rough estimate of the amplitude *A* of the oscillation, including uncertainty. [Express your final answer in cm, and express your uncertainty to one significant figure.]

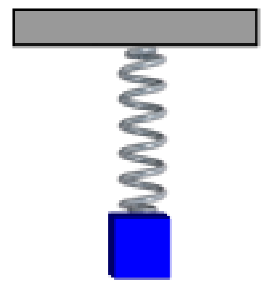

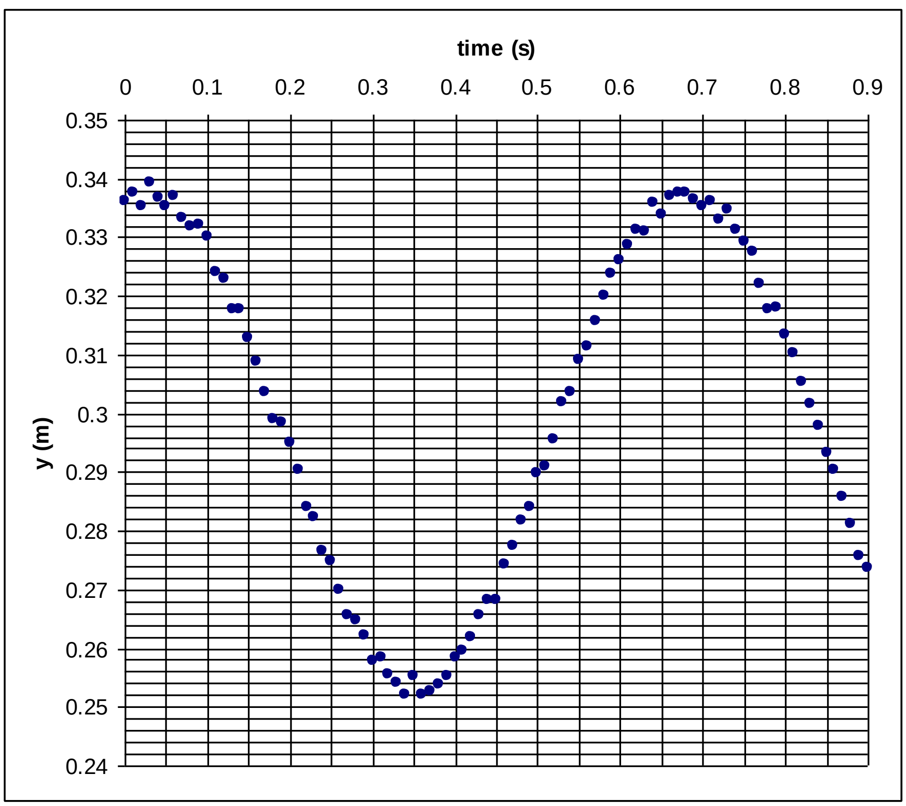

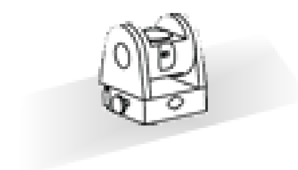

12. In the Practicals a student uses the computer-based ultrasonic motion sensor to measure distance versus time for a cart. The cart has low-friction wheels and is on a straight 2.2-metre aluminum track that is tilted at some angle. The cart is released from rest at time $t = 2.7$ s, and distance data is collected at a rate of about 10 measurements per second until $t = 9.1$ s. Over this time, the distance increases from 57 cm up to 200 cm as the cart slowly rolls down the slope. A polynomial function of t is fit to the data:

$$\text{Distance (cm)} = a0 + a1\ t + a2\ t^2$$

Where t is the time in seconds, and a0, a1 and a2 are the fit coefficients. The best fit gives:

- a0 = 64 ± 0.31 cm
- a1 = −10.95 ± 0.11 cm/s
- a2 = 2.8792 ± 0.0094 cm/s^2

From this fit, what would you conclude is the acceleration of the cart?

13. In the figure the frictionless Track is at an angle θ with the horizontal. The Cart has a mass M, and is connected to a hanging mass $m = 0.075 \pm 0.002$ kg by a massless string over a massless, frictionless pulley.
For an angle $\theta = 5.2 \pm 0.1°$ the masses are in equilibrium, i.e. if they are at rest they remain at rest and if they are moving at some speed they continue moving at that speed. Calculate the value of the mass of the cart M. Express your final answer in kg. [You may use the small angle approximation which is that $\sin\theta \approx \tan\theta \approx \theta$ when θ is measured in radians.]

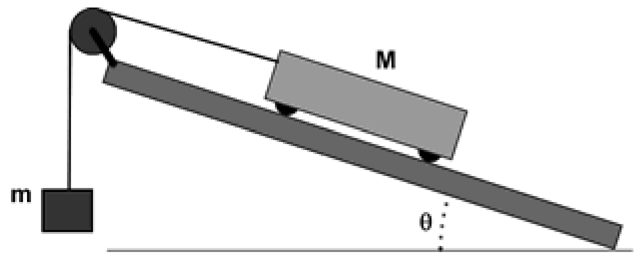

14. A small scale is used to measure the individual mass of N different pieces of gravel. The estimated mean of the N measurements is 20 g and the uncertainty of the mean is 4 g. Next the mass of a pile of the gravel is found to be 6000 ± 20 g. What is the best estimate of the number of pieces of gravel in the pile?

15. You wish to calculate the density of a solid sphere of foam by measuring its mass and diameter. Density = mass/volume, and the volume of a sphere = $\frac{\pi}{6}d^3$, where d is the diameter. You can measure the mass of the sphere of foam to a percentage uncertainty of 1%, and the diameter to a percentage uncertainty of 2%. What will be the percentage uncertainty in the density of foam that you compute?

16. You are counting the number of worms per bucket of dirt in a farmer's field. After counting the worms in N different buckets of dirt, you compute the estimated mean, $\bar{x}$, and estimated standard deviation, σ, of the numbers. If you continue your measurements until you have counted the numbers of worms in 100N buckets, what would you expect to be the standard deviation of the numbers?

17. A research student on a farm counts the number of worms per bucket of dirt in a farmer's field. She counts the worms in six different equally-sized buckets, as recorded in the table below.

Number of worms:	70	80	89	68	103	71

i. Based on these data, what is the mean number of worms per bucket? Be sure to include the uncertainty of the mean, and express both the mean and the uncertainty to the correct number of significant figures.
ii. How many of the six measurements lie within plus-or-minus one standard deviation of the mean?