Uncertainty Analysis Workshop

When: Begin: January 17 ; Due: January 27

Read: Taylor, *An Introduction To Error Analysis*
Chapters: 1, 2, 3, 4, 8 (Now); Read entire book over the semester.

`http://www.physics.csbsju.edu/stats/`
pages: Mean, Standard Deviation, etc, Ordinary Least Squares (Regression)
Beyond Ordinary Least Squares, WAPP: Fit to data with y-errors
ANOVA: ANalysis Of VAriance between groups

Homework: (work in lab notebook) 3.18, 3.26, 3.36, 3.48, 3.50, 4.5, 4.18, 4.28, 8.26 (see following)
8.26: Do not take the log of the data, instead enter the count data into *WAPP*, use the proper error and functional form and make a semilog plot. Calculate τ *and its error* from the fit.
Read Web: Ordinary Least Squares (Regression) and record tree-age data.
Read Web: Beyond Ordinary Least Squares and *read* problem 8.17.
Use above web site to fit and display several types of fits to the tree-age data. Tape in your notebook a plot of several of these fit lines. Label the lines with the type of fit. Select what is in your opinion the best fit. Explain *why* you think it is the best fit. Make a table displaying the parameters from the various fits. Circle the minimum and maximum slope and y-intercept.

Lab: Complete experiments 1 and 3 below. (Many of you did 2 last semester.)
Record the results in your notebook.

References: Bevington, *Data Reduction and Error Analysis for the Physical Sciences*
Press *et al.*, *Numerical Recipes*

Abstract

These three experiments are intended to demonstrate uncertainties as they can be found in real experiments and to introduce you to more sophisticated measures of these uncertainties. The first experiment deals with the problem of definition (p. 46) which occurs when the phenomenon to be measured is "fuzzier" than the device used to measure it. The second experiment deals with random event counting in which the count itself provides the uncertainty. The final experiment introduces statistical uncertainty.

1 Finding the Focal Length of a Lens

1.1 Purpose

The purpose of this experiment is to introduce the problem of definition. In many of our labs the problem of definition is the primary source of uncertainty.

1.2 Theory

(See Halliday, Resnick & Walker chapter 35, particularly section 35.6.) The light leaving an illuminated object (like an illuminated slide) can be collected by a lens and used to produce a real image on a card located on the other side of the lens. If o is the distance between the slide (the "object") and the lens, and i is the distance between the card (with focused image) and the lens, then it can be shown that

$$\frac{1}{i} + \frac{1}{o} = \frac{1}{f}$$

where f (a constant) is the focal length of the lens. You are to try to verify this equation and determine f. Note that, by this equation, o must be greater than f if a real image forms (i.e., $i > 0$).

1.3 Procedure

Move the holder with the slide to the 0 cm mark on the optical bench. Set-up the lamp and shade to illuminate the slide, sending the light toward the lens. Start with the lens at the 19 cm mark and the card near the 100 cm mark. Align the lamp and move the *card* so that the image on the card is as clear as possible. Note that the quality of being "focused" is hard to define, i.e., a range of measurably different card positions all look focused. This is the problem of definition. Record half the range of possible card positions as the uncertainty in the card position. Have your lab partner do the same and see how your ranges and best values compare. (Note that as a result of the large uncertainty in the card position, i (position of the card minus position of the lens) will have a large uncertainty, whereas o (position of the lens minus position of the slide) will have a small uncertainty. Thus, following the standard practice of putting the small-error quantity on the x-axis, o will be the independent variable and i will be the dependent variable.) Repeat this process producing at least eight well-chosen lens positions. (Note that *well-chosen* means that you are not essentially repeating an earlier measurement, i.e., make sure that either i or o has changed substantially between your measurements. Both i and o should span the range 20 to 50 cm.)

As you are doing the experiment, fill in a table with slide position (which might as well be fixed), lens position, card position, o, and i. (Of course each quantity has an uncertainty.) You should find that the uncertainty in i is not constant. Use the appropriate mode of *WAPP* for y-errors and chose the proper functional form for the expected relationship. Find f from the fit parameters.

1.4 Using *WAPP* on the Web

Start your favorite web-browser. (I use *Netscape*.) Enter the web address:

`http://www.physics.csbsju.edu/`

Select Statistics and then *WAPP*: Fit to data with y-errors Note that your error in $y = i$ varies; neither absolute nor percentage errors are appropriate. You must: "Enter separate y errors for each data point". The program will then allow you to set the error for each data point. Select the appropriate functional form. The program will print out your results. *DO NOT* play the error game of re-guessing your original error to get what somebody told you is the "right" χ^2. Record the parameter values and errors. Tape a copy of the computer results into your notebook. You will need them to calculate f!

You should also produce a plot (to appear in your notebook). Plots will be best reproduced on postscript printers. Select the "PDF" format; the program *Acrobat* should launch and allow you to see and print the plot.

As you know, many functional forms can be linearized by transforming variables before you plot. For example, plotting $\log(x)$ and $\log(y)$ will straighten out a power-law function. Log-log paper makes the transformation automatic by labeling lines with x while spacing the lines according to $\log(x)$. In this experiment plotting $\frac{1}{x}$ and $\frac{1}{y}$ should produce a straight line; we need "inverse-inverse" paper to make the transformation automatic. Such odd sorts of paper are available on *WAPP* by simply selecting the proper scale options. Produce such a plot and tape it in your notebook.

1.5 Results

Report your value for the focal length of the lens with an uncertainty. Include both linear and inverse-inverse plots.

2 Uncertainty in Random Event Counting

2.1 Purpose

The purpose of this experiment is to notice that in counting random events, like nuclear decays, for a fixed time, repeated measurement of the count yields different values. The randomness of nuclear decay necessarily produces uncertainty in the original count. You are to demonstrate that the standard deviation of the number of counts during a fixed time period (60 seconds) is about equal to the square root of the average number of counts. Thus in future labs you can estimate the uncertainty in one count value without repeating the count; just use the square root of the count as the uncertainty in the count.

2.2 Theory

The mathematics behind the physics is the Poisson distribution: see Taylor, chapter 11 or Bevington, section 3–2.

Nuclear decays frequently produce energetic particles called alpha (α) particles. (An α particle consists of two protons and two neutrons bound together and is exactly the same thing as a helium (^{4}He) nucleus.) This experiment seeks to count these sub-atomic particles, as they pass through a screen. The screen is doped Zinc Sulfide, ZnS(Ag). When an α particle strikes the screen it excites some of the electrons in the material to a higher energy state. When one of the excited electrons returns to its normal state, a photon of light is emitted. The emitted photons then enter the photomultiplier tube and hit the cathode. The cathode is made from a photoelectric material which emits a fixed number of electrons per photon hit (i.e., the incoming photon is converted to several electrons). These electrons are then accelerated by a high voltage and slammed into another electrode; 3–6 new electrons are ejected per electron hit (i.e., the electrons are *multiplied* and the current amplified). These newly ejected electrons are again accelerated and directed onto the next electrode where again they eject 3–6 times as many electrons. This happens ten to twelve times. The resulting current pulse from the anode of the photomultiplier is large enough to be detected and counted as an event. *DANGER*: high voltage is required to accelerate the electrons; the definition of **high voltage** is a voltage large enough to kill a person. Let's not test to see if the definition is accurate!

2.3 Procedure

CAUTION: The radioactive α source we are using is contained within the black photomultiplier tube. Do not take it apart to see this source.

1. Ludlum 2000 scaler should have been left on by the previous student. If not turn the main power switch to *LINE* and let it warm up for about 10 minutes.

2. All other knobs have been preset, but make sure the high voltage (HV) is turned to about 4.5 (so the meter reads ∼650 V), the minutes are set to 1, and the multiplier is set to X1. If you think something is not set right, *ask Tom or Dan* to check it out. Do not worry if there are numbers on the display. These will be cleared automatically when you press the *COUNT* button.

3. Press the *COUNT* button. The red light should go on and you should see the Ludlum display go to zero and then start counting up as it records events. It will stop automatically after it has counted for a minute. The red light will go dark when the count-time has expired.

4. Record the count.

5. Go to 3 and repeat until you have ten counts at one minute each.

2.4 Results

Report the average and standard deviation of the count over 10 trials. (Either learn to use the stats functions on your calculator or use the web: Mean, Standard Deviation, etc.) Calculate the square root of the average count and compare it to the standard deviation. One is supposed to find $\sqrt{N} \simeq \sigma_N$.

3 Statistical Uncertainty

3.1 Purpose

The purpose of this experiment is to investigate a more sophisticated measurement of uncertainty: standard deviation and standard deviation of the mean (SDOM). In this experiment we use a die to simulate a fluctuating measurement (e.g., the fluctuating least significant digit on a digital meter). In this case, what should be recorded as the measurement and what should be recorded as the uncertainty? (The answers are discussed in Taylor; in short: the mean and the standard deviation of the mean.) Clearly for an individual measurement the range of fluctuation (i.e., the standard deviation) indicates the uncertainty. However in finding the range of fluctuation, more data is collected—often the average (mean) of that data better represents the measurement than any single measurement. But if we record the mean as the value, what is the uncertainty in that value? Surely averaging helps reduce the fluctuation—that is if we were to repeat the whole process of finding the mean again and again the means would fluctuate less than the values averaged. Taylor reports (Equation (4.14)) a formula that predicts the fluctuation in the means (SDOM) from the fluctuation in the values and the number of values averaged. If this formula is correct, in future labs you can estimate the uncertainty in your mean just using the values you recorded to calculate that mean.

3.2 Theory

(See Taylor chapter 4 or Bevington chapter 2.)

3.3 Procedure

Find five six-sided dice. Roll the five dice 16 times, recording the value on each die (e.g., the red die value, the green die value, etc.) and their sum in a table.

3.4 Results

Use a calculator or the above statistics web site to automatically calculate the mean and standard deviation of each of the six columns of data (five dice and sum). For each column, report the percentage of the values that lie one standard deviation or less from the mean; no "rounding"! (cf., middle p. 101). See if Equation (3.16) is approximately satisfied for your sum. (This means calculate *both* sides of Equation (3.16) and report how accurate the supposed equality actually is.) How well does Equation (3.4) work? Calculate the standard deviation of your five means. You should find that the means have a smaller standard deviation (uncertainty) than the raw measurements, i.e., the fluctuation (error) in an average is less than the fluctuation of the numbers that produce that average. In fact, Equation (4.14) claims

$$\sigma_{\bar{x}} \simeq \frac{\sigma_x}{\sqrt{N}}$$

Check it out! (This means calculate *both* sides of Equation (4.14) and report how accurate the supposed equality actually is.)

A "loaded" die is one modified so it produces "unusual" results. (For example, producing on average more 6s than 1s, rather than the expected unbiased outcomes.) A common way to detect unusual distributions is ANOVA. Using the ANOVA web page (with "copy & paste data entry") enter the results for your five dice. Hardcopy the results and produce a plot displaying the group means with 95% confidence intervals. Tape these hardcopies into your notebook. Did you detect a loaded die? If so, which one is loaded and what is the evidence that it is loaded?

4 Summary

The way to estimate statistical uncertainties is to repeat the experiment many times and to report the standard deviation of the mean as the statistical uncertainty. However, we instructors have a hard time convincing you to do the same experiment a few thousand times, so we have often come up with "short cuts", like trying to guess how the results would differ if you actually did the experiment again or using the manufacturer's specifications. These short cuts are really wrong. As a result you will often miss-guess your uncertainties by a factor or 2 or more, and thus produce reduced χ^2s that deviate from 1 by a factor of 4 or more. We instructors know this—there is no point in trying to hide what somebody told you was a "bad" χ^2. A not uncommon (but not really legitimate) approach is to make a series of measurements that you think should to fall on a "smooth curve", and then use the deviation of those measurements from that "smooth curve" as a measure of your uncertainty. (The legitimacy of this approach is much reduced if the "smooth curve" is really a fit curve, but in fact this is how our elementary fitting programs work.) About the only case where there is a fully legitimate short cut in finding the uncertainty is in counting random events (and then only if the events are known to be truly random and if the apparatus is known to be working properly—and generally the only way to be sure of all that is to repeat the experiment). I must finally add that the magic known as "resampling" seems to provide a "free lunch"— one data-set may play the role of many. So if at some point you must estimate errors but cannot repeat the experiment, investigate resampling.

Statistical uncertainties are the easy part; systematic uncertainties are so tricky this workshop has been largely silent about them. (The summary statement for systematic errors is do the measurement again a totally different way. But in real life that job is often given to another experimenter and you are left to make an *informed* guess about calibration errors.)