PHYS 310 – HOMEWORK #4 SOLUTIONS

Reading:

- Hughes & Hase, Chapters 5, 6.1-6.2

Problems Due Tues 16 February

There are three problems in this set. You must download the data for these problems from the website http://www.eg.bucknell.edu/physics/ph310/fit1.html or Moodle.

1) Use the "hw4-1.dat" file for this problem. These data come from twenty (20) experiments nominally measuring the same physical quantity. The data for each measurement (point) is on a single line in the file. The first number on each line is the value of the measured quantity, and the second is the uncertainty in the measurement.

   (a) What value do you quote for this quantity based on the data? (Include an uncertainty.)

In ipython:

```
y, u = loadtxt('hw4-1.dat', unpack=True)
w = 1/u**2
wmean = sum(y*w)/sum(w)
wunc = 1./sqrt(sum(w))
```

I get 14.469 ± 0.059

   (b) Each data point has its own uncertainty σ_i. How many of the data points lie within $1\sigma_i$ of the mean value you determined?

Counting ± 1σ points is easy on a plot. From inspection of the plot below, it looks like 14 / 20 = 70% are within $1\sigma_i$ of the mean.

```
clf()
x = arange(len(y))
errorbar(x, y-wmean, u, fmt='ko')
xlim(-1, 21)
plot((-1,21),(0,0), 'r-')
xlabel('Data Point Number')
ylabel('Data - Weighted Mean')
```

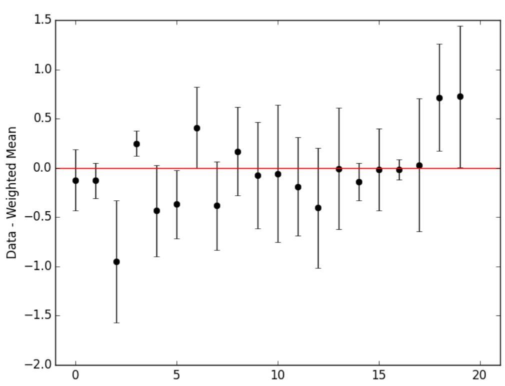

We can also get python to do the counting for us. Calculate the normalized residuals, and count the number with magnitude < 1.

```
z = (y - wmean) / u
sum(abs(z) < 1)
```

Ok, so it was 15. My eyes aren't what they used to be.

(c) What is the goodness-of-fit parameter χ^2 for these data? The definition of χ^2 is given in Eq. (5.9) of Hughes & Hase, and we calculated χ^2 for the cases of fitting to a linear and a quadratic in class. When we are just determining the mean, the function *y(x)* in Eq. (5.9) is simply a constant, i.e., the value of the mean that you determined.

Since I've already defined the normalized residual above, I just need to sum the squares.

```
chisq = sum(z*z)
```

The result is 14.7.

2) Use the "hw4-2.dat" file for this problem. The data come from an experiment in which there is a suspected linear relationship between measured values of x and y. The data for each point is on a single line in the file. The first number on each line is the value of x, the second is the value of y, and the third is the uncertainty of y. Uncertainties in x are negligible.

(a) Perform a weighted linear fit to these data.

I simply modified the "goodfit.py" script to read in 'hw4-2.dat' instead.

(b) Plot your normalized residuals. How many of the data points lie within 1σ of the line you determined?

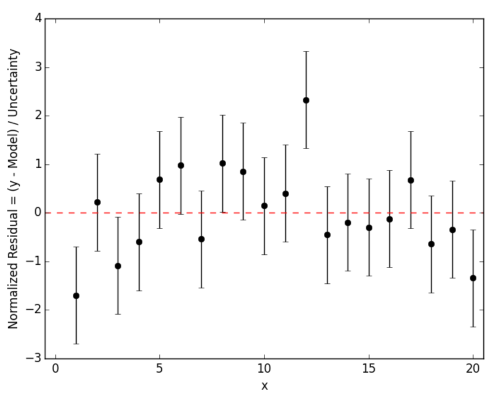

I count about 15 / 20 = 0.75 of the data to be within 1σ of the line.

(c) Use Eq. (5.9) from Hughes & Hase to calculate the goodness-of-fit parameter χ^2 for these data.

Happily, the program does this for you. I find χ^2 is about 16.7.

(d) Does a linear fit to the data appear to be reasonable?

The fit looks reasonable owing to the results of parts (c) and (d). On the other hand, there remains a curvature in the residuals that's worrying. I suspect a second order fit would be better, and the uncertainties were overestimated.

(e) What value do you quote for the slope and intercept based on the data? (Include uncertainties.)

Polynomial coefficients
slope = 2.63 ± 0.15
intercept = 1.8 ± 1.9

3) Use the file "hw4-3.dat" on the website for this problem. The data come from an experiment in which the relationship between x and y is suspected to be

$$y = a_1 \sin(2\pi x) + a_2 \sin(4\pi x),$$

where a_1 and a_2 are the parameters to be determined. The data for each point is on a single line in the file. The first number on each line is the value of x, the second is the value of y, and the third is the uncertainty in y. Uncertainties in x are negligible.

(a) Perform a linear fit using the assumed functional form.

I modified 'mlr.py' to perform the fit. The main modification is to redefine x1 and x2, and to remove x3.

```
x1 = sin(2*pi*x)
x2 = sin(4*pi*x)
…
A, model, covariance = LinearModelFit(zip(x1, x2), y, u)
```

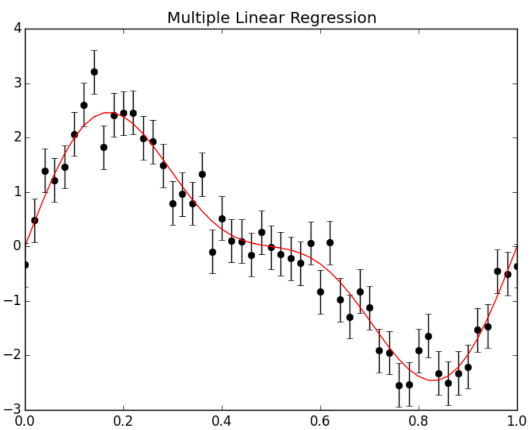

(b) Plot your residuals.

In python:

```
clf()
errorbar(x, y-model, u, fmt='ko')
xlabel ('x')
ylabel('residual = y - model')
plot((-0.1,1.1), (0,0), 'r-')
xlim(-0.1, 1.1)
savefig('hw4-residual.png')
```

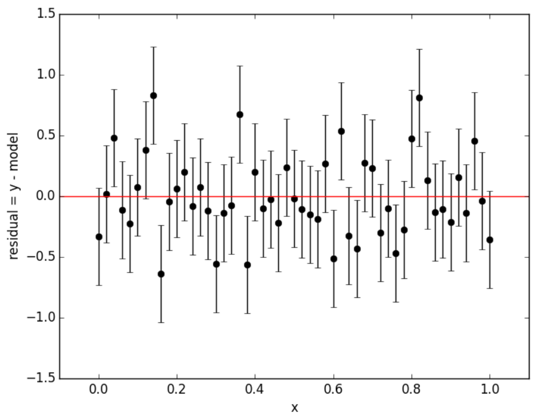

(c) Use Eq. (5.9) from Hughes & Hase to calculate the goodness-of-fit parameter χ^2 for these data.

In python:

```
z = (y – model)/u
print sum(z*z)
```

I find χ^2 is about 36.6.

(d) Is your fit good?

Judging from the residuals and χ^2, this fit is good. In fact, it's too good because χ^2 < the number of data points. Somebody overestimated the uncertainty.

(e) Give your values for a_1 and a_2 (including uncertainties).

The code reports
$a_1 = 1.965 \pm 0.069$
$a_2 = 0.879 \pm 0.069$