Homework #3 solutions, AST 303, Fall 2008

100 points total

General grading rules: 3 points off per arithmetic, algebraic, or conceptual mistake. 1 point off for grossly too many significant figures. 1 point off for giving answers without units.

1. More IDL practice (*35 points*)

a. (*10 points*) Make a color version of the WMAP-measured temperature fluctuations, and print out the map. *Hint: Try using the* `atv` *command to display the image that is output by the* `healpix_to_image` *command; you may have to run it twice to get it in full color.*
Full credit for any reasonable image of the map. As many of them will be using a black-and-white printer, a black-and-white image is acceptable.

b. (*5 points*) Make a histogram of the temperature distribution.
Take a point off for axes that are not labelled. Take another point off if there is no written description of how the binning was done.

Note that the output of the `fits_read_map` command is in milliKelvin (as can be confirmed by reading the header of the original fits file). The output of the `healpix_to_image` command is on a scale from 0 to 255, for display; it is no longer in physical units. So you want to use the former for all calculations here. The histogram you get is nicely Gaussian.

c. (*10 points*) Compute the mean, variance, skewness and kurtosis of the map.
Using canned routines in IDL for calculating this is a bit missing the point here, but is acceptable. There is more than one way to define skewness and kurtosis (do you normalize by σ^3, σ^4 or not?); full credit for doing it either way, as long as it is explicit which is done. If it is not explicit, take 3 points off.

Few people were explicit about this; the IDL routine `moment` does the work, and it does renormalize things. Thus the skewness of the quantity x that IDL uses:

$$\frac{\left\langle (x - \langle x \rangle)^3 \right\rangle}{\sigma^3}$$

and the kurtosis is:

$$\frac{\left\langle (x - \langle x \rangle)^4 \right\rangle}{\sigma^4} - 3,$$

where

$$\sigma^2 = \left\langle (x - \langle x \rangle)^2 \right\rangle$$

is the variance. By the way, do you understand why 3 is subtracted from the kurtosis? Ask yourself what the fourth moment for a Gaussian function is, and come talk to me if this isn't clear.

When I do this for the WMAP data, I get:

$$\text{mean} = 1.47 \times 10^{-7}\,\text{mK}$$

It makes sense that this is tiny; the mean has already been subtracted off this fluctuation map.

$$\text{variance} = 0.005\,\text{mK}^2$$

Note the units. This corresponds to a standard deviation of 0.071 mK, or about 2×10^{-5} of the mean temperature of 2.7 K.

$$\text{skewness} = -0.116$$

This is unitless, as described above. The other way of defining skewness doesn't normalize by σ^3; you get $4.13 \times 10^{-5}\,\text{mK}^3$.

For the kurtosis, one gets the dimensionless number 0.351; if one doesn't normalize, one gets $8.9 \times 10^{-6}\,\text{mK}^4$.

It turns out that there is a great deal of interest in the skewness and kurtosis of this map. If the distribution is purely Gaussian (as is predicted in simple models of early universe inflation), the skewness and kurtosis should be accurately zero; deviations from zero may tell us about processes in the early universe. To interpret the numbers we've just calculated, we'd need to think about the statistical uncertainty on them, and ask whether they were consistent with zero. Lots of people have looked in detail at this question, and in fact, there is a raging controversy at this point about whether there is evidence for non-Gaussianity in the maps. David Spergel, who is a member of the WMAP team, says that there is no such evidence.

d. (*10 points*) In IDL, try typing the following:

```
print,2^14
print,2^15
print,2^16
print,7*1000
print,7*10000
print,7*10000*2
print,7*10000*2L
print,7*10000L*2
```

What happened? Why? IDL uses what are called 16-bit integers by default; with this in mind please explain quantitatively how each of these results came about.

Solution:

2^14 gives 16384; this is correct.

2^15 gives −32768. This is because we are limited to 16-bit numbers. One of the bits is taken up by the sign (+ or −), and one is the zero-order bit. Thus the largest integer that IDL can render correctly is $2^{15} - 1 = 32767$; when another bit is added, it incorrectly overflows to the sign.

2^16 gives 0. This would (correctly) be written as a 17-digit number in binary: 10000000000000000, but the initial 1 is not renderable, leaving 0.
7*1000 gives 7000; that's correct.
7*10000 = 70000 = 65536 + 4464; and 65536 = 2^{16}, which we've seen is rendered as zero. So this gives 4464.
7*10000*2. The calculation of 7*10000 is done first, giving 4464; doubling it gives 8928.
7*10000*2L. The 'L', standing for ''long', says to use 32-bit integer arithmetic. However, the first part of the calculation is done still as a 16-bit calculation, giving the same result as above.
7*10000L*2. Now the calculation is done fully with long integers, giving the correct answer of 140,000.

2. The relationship between broad-band and monochromatic magnitudes (*25 points*)

One typically observes a star through a broad filter (as we'll see in the next problem), while the magnitude scale refers to a monochromatic flux. Consider a star whose spectrum is described by a power law:

$$f_\lambda = A\left(\frac{\lambda}{\lambda_0}\right)^\gamma.$$

You measure this star with a photon-counting detector through a filter which is a tophat: i.e., it lets 100% of the light through in the wavelength range $\lambda_0 - \Delta\lambda$ to $\lambda_0 + \Delta\lambda$, and no light through elsewhere. Calculate the flux of this star measured through the filter as a function of the power-law index γ, and expand in orders of $\Delta\lambda$, to the first non-vanishing term in $\Delta\lambda$. While the flux density at λ_0 is independent of γ, the flux through the filter does depend on γ. A typical range of γ is from -2 to $+2$; by what factor does the flux through the filter vary over this range? Do this calculation for $\Delta\lambda/\lambda_0 = 0.02$ (an intermediate-band filter) and $\Delta\lambda/\lambda_0 = 0.2$ (a broad-band filter).

Solution: The flux through a filter R is:

$$\frac{f = \int f(\lambda)R(\lambda)\,\lambda\,d\lambda}{\int R(\lambda)\lambda\,d\lambda}$$

(Note that we can write this in terms of $f(\lambda)$ or $f(\nu)$; in the latter case $d\lambda/\lambda \to d\nu/\nu$). For the tophat filter and the spectrum we have above, this becomes, simply:

$$f = A\frac{\int_{\lambda_0-\Delta\lambda}^{\lambda_0+\Delta\lambda} \lambda\,d\lambda \left(\frac{\lambda}{\lambda_0}\right)^\gamma}{\int_{\lambda_0-\Delta\lambda}^{\lambda_0+\Delta\lambda} \lambda\,d\lambda}.$$

To make life easier, I am going to make the substitution, $x \equiv \frac{\lambda}{\lambda_0}$, and $\delta \equiv \frac{\Delta\lambda}{\lambda_0}$. Then

$$f = A\frac{\int_{1-\delta}^{1+\delta} dx\,x^{\gamma+1}}{\int_{1-\delta}^{1+\delta} x\,dx} = \frac{A}{\gamma+2}\frac{(1+\delta)^{\gamma+2} - (1-\delta)^{\gamma+2}}{2\,\delta}.$$

In the last step, we assumed that $\gamma \neq -2$; for $\gamma = -2$, numerator and denominator are identical, and we get $f = A$ exactly.

Let's expand this expression in δ, which is usually much less than one (and certainly true in our case!). We find we have to go to third (!) order to see a non-vanishing term in δ; we get, after a bit of algebra:

$$f = A\left(1 + \frac{(\gamma+1)\gamma}{6}\delta^2\right).$$

So we now have what we want, the fractional change in the received flux as a function of the shape of the spectrum (as parameterized by γ) and the width of the filter (δ). We have already seen that this change is zero for $\gamma = -2$; we can see that it is zero as well (to the level of our expansion) for $\gamma = -1$ and $\gamma = 0$. For $\delta \equiv \Delta\lambda/\lambda = 0.02$, the correction is of order $7 \times 10^{-5}\gamma(\gamma+1)$. That is indeed a tiny correction, less than 0.01% for most values of γ. For $\delta = 0.2$, we get $7 \times 10^{-3}\gamma(\gamma+1)$, i.e., roughly a 1% effect, which is difficult, albeit not impossible to measure.

20 points for the calculation of the flux algebraically, although take 10 points off if there is no expansion in powers of δ. 5 points for any reasonable plugging in of numbers.

3. The magnitude of an A star (*40 points*)

On the course home page, you will find the tabulated spectrum of an A0V star; the first two columns are wavelength (in Ångstroms) and flux density f_λ (in units of $10^{-17}\,\mathrm{erg\,s^{-1}\,cm^{-2}\,Å^{-1}}$; this spectrum is from a compilation by Pickles 1998, PASP, 110, 863).

a. (*15 points*) Calculate the apparent magnitudes of this star, on the AB system (see Homework 1, Problem 2), through the SDSS filters u, g, r, i, and z. Given that the AB system makes reference to a monochromatic flux, you can take the effective frequency of each filter to be the first moment of the filter bandpass (i.e., the mean of the frequency over the bandpass), which you will have to calculate. You can find the filter curves at: `http://www.sdss.org/dr5/instruments/imager/filters/index.html` (you want to use columns 1 and 2). In carrying out the integration over the filter curves, you can turn the integrals into sums. But you will have to deal with the fact that the star spectrum and the filter curves are defined on different grids of wavelength. You should interpolate the filter curve to the points on which the star spectrum is defined; this can be done with linear interpolation, or using so-called spline routines which you may be able to find in IDL. All calculations should be done to four significant figures.

Solution: To calculate the effective flux density through the filter, we integrate over the spectrum and filter, using the expressions we saw above:

$$f = \frac{\int f(\lambda)R(\lambda)\,\lambda\,d\lambda}{\int R(\lambda)\,\lambda\,d\lambda}$$

where f is the spectrum of the star. This of course will give us a result in units of $10^{-17}\,\mathrm{erg\,cm^{-2}\,s^{-1}\,Å}$, and we need to convert to $\mathrm{erg\,cm^{-2}\,s^{-1}\,Hz^{-1}}$ to calculate magnitudes. To do this, we remember that the integrated flux is the same whether we calculate with respect to wavelength or flux;

$$\int f(\lambda)d\lambda = \int f(\nu)d\nu$$

Differentiating both sides tells us:

$$f(\nu) = f(\lambda)\frac{d\lambda}{d\nu} = f(\lambda)\frac{c}{\nu^2}$$

(where the dropped minus sign is immaterial). Following the units, and remembering that $c = 3 \times 10^{18}$Å/sec, we get:

$$f(\nu) = f(\lambda)(10^{-17}\,\mathrm{erg\,cm^{-2}\,s^{-1}\,Å}) \times \left(\frac{3 \times 10^{18}\,\mathrm{Å\,sec^{-1}}}{(\nu/10^{14}\,\mathrm{Hz})^2 \times 10^{28}}\right) \times \frac{1\,\mathrm{Jy}}{10^{-23}\,\mathrm{erg\,cm^{-2}\,s^{-1}\,Hz}}.$$

$$f(\nu) = 3 \times 10^{-27}\, f(\lambda)(10^{-17}\,\mathrm{erg\,cm^{-2}\,s^{-1}\,Å})\frac{1}{(\nu/10^{14}\,\mathrm{Hz})^2}.$$

Here I scaled ν to the nearest order of magnitude.

To do the integral over wavelength, I notice that the spacing on which I have the spectrum, and that for the filter, are quite different. I can simply linearly interpolate the filter for the missing pieces. In my code, I did something more sophisticated called a spline, which we will find ourselves discussing in a few lectures.

The magnitudes are given by $-48.6 - 2.5 \log f(\nu)$, calculated at the effective wavelength. The expression above for the integrated flux is normalized such that it ends up in units of erg/s/cm^2/Å; following the above, it is easy to convert to units of $f(\nu)$, or erg/s/cm^2/Hz.

Here are the results:

Filter	*mag*	m_2
u	22.24	22.42
g	21.27	21.19
r	21.50	21.48
i	21.68	21.67
z	21.84	21.71

Here, m_1 is the magnitude calculated through the filter, while m_2 is the magnitude calculated at the effective frequency of the filter (see part b).
Some will calculate the effective frequency of the filter in this part, some in the next part; fine either way. Depending on whether one is counting photons or energy, the integrals used in the flux calculation will either have or not have that extra factor of λ. *Either is fine (and it makes a difference of a few hundredths at most in the final magnitude calculation), as long as it is clearly explained which is used. If this isn't clear, take off 4 points.*

Note that these are not the same thing! Indeed, the attached figure, which shows the spectrum, the filters, and the fluxes calculated the two different ways (blue is doing the integral properly, while red simply determines the flux density at the effective wavelength), demonstrates that the biggest difference occurs in the u band, where the spectrum has a huge jump (the so-called Balmer break).

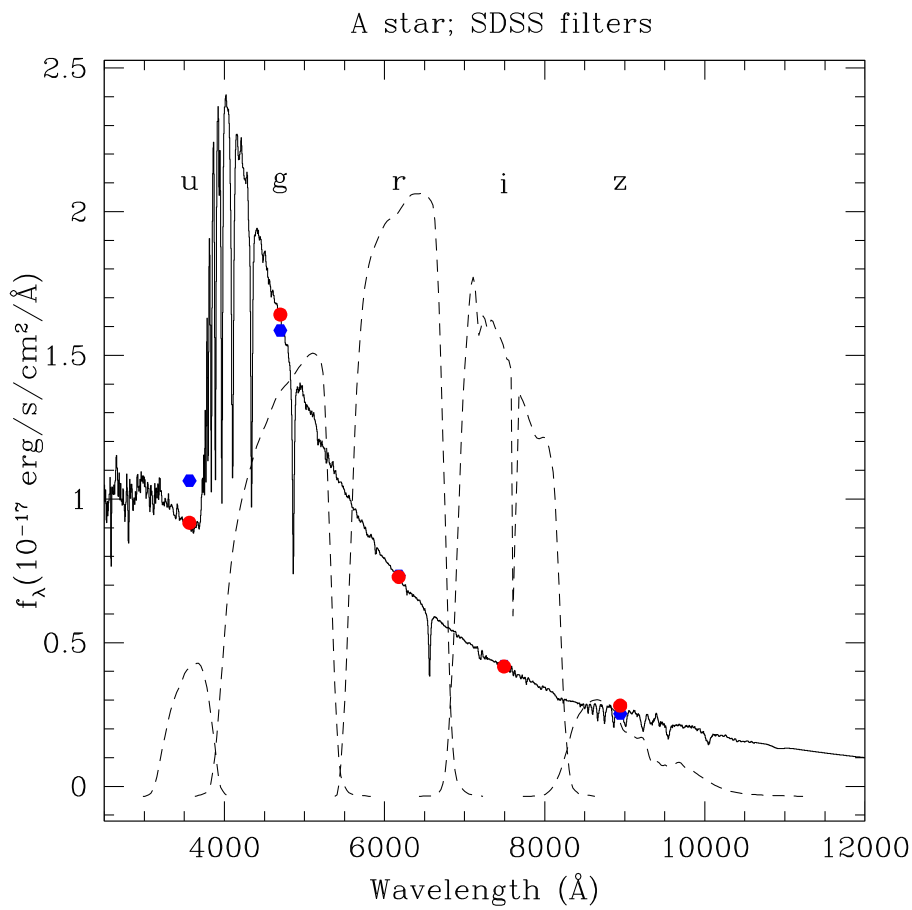

Figure 1: Figure for Problem 2: the spectrum of A star, the SDSS filter curves, and the mean flux density in each filter, calculated two different ways.

b. (*10 points*) Now do the calculation the naive way, namely calculate the magnitude of the star using the monochromatic flux at the effective frequency of each filter. How much does your answer differ from that which you calculated directly? You should have found the discrepancy between the two is worst in u; why is that?

Solution: The real work here is to calculate the effective frequency of each filter. It is given by:

$$\nu_{eff} = \frac{\int d\lambda\, \nu\, R(\lambda)}{\int d\lambda\, R(\lambda)}.$$

It is actually ambiguous whether the integrand should be $d\lambda$, as I've done here, or $\lambda d\lambda$, which gives a photon-weighted effective wavelength. In order to do these integrals, I note that the filter curves are defined every $\Delta\lambda = 25$ Å, so it is easy to turn the integrals into sums:

$$\int d\lambda\, \nu\, R(\lambda) \rightarrow \sum \Delta\, \lambda\, \nu R(\lambda).$$

(We will discuss later in the course more sophisticated things we can do for the calculation of the integral). What I will do is calculate the sum using the effective wavelength at the *midpoint* of each bin. The code I wrote checks the width of each bin explicitly, rather than assuming that they are all spaced by 25Å (although they are, in fact). As long as I was at it, I calculated the effective wavelength as well (which is *not* just c/ν_{eff}). To 4 significant figures, I get:

Filter	ν_{eff}	λ_{eff}
u	8.418×10^{14} Hz	3574Å
g	6.384×10^{14} Hz	4731Å
r	4.856×10^{14} Hz	6198Å
i	4.004×10^{14} Hz	7512Å
z	3.354×10^{14} Hz	8974Å

Given these effective frequencies, it is straightforward to calculate the flux at those frequencies; we tabulated them in the table above.

7 points for calculation of effective frequency (using either of the weighting functions), given either here or as part of part (a), and 3 points for calculation of magnitudes.

c. (*5 points*) On a Vega-based magnitude system, the magnitudes of an A0V star are independent of filter, by definition. Is this the case for the AB system? Why or why not?

Solution: No, this is not the case for the AB system, simply because constant AB magnitude with wavelength would correspond to a constant $f(\nu)$, with the spectrum clearly does not follow.

d. (*10 points*) I asked the calculations to be done to four significant figures, i.e., magnitudes to the nearest 0.01 mag. What percentage change in flux does 0.01 mag correspond to? To measure a magnitude to that accuracy, what is the minimum signal-to-noise ratio that your measurement needs to be made to? If the signal-to-noise ratio is set by Poisson statistics of the incoming photons alone, how many photons do you need to detect from the star?

Solution: The flux density f of an object is related to its magnitude m by:

$$m = -48.60 - 2.5 \log f,$$

or

$$f \propto 10^{-0.4m}.$$

Therefore a change of 0.01 mag corresponds to a fractional flux change of:

$$\frac{\delta f}{f} = 10^{0.4\times 0.01} - 1.$$

To get a sense of this, I convert to base e. As $\ln 10 = 2.303$ and $2.303 \times 0.4 = 0.921$, this gives

$$\frac{\delta f}{f} = e^{0.009} - 1 \approx 0.009,$$

or just under 1%. Thus to a rough approximation, for small magnitude differences, a change of x magnitudes corresponds to a fractional change of x.

The signal-to-noise ratio is just the ratio of the flux to the error, i.e., roughly 100.

The signal to noise ratio for a Poisson process with N photons is $\sqrt{N}$, as we saw in class. So with a signal-to-noise of 100 requires 10^4 photons.

4 points for the calculation of the change in flux, and three points each for the S/N and the number of photons. Doing the calculation to another two decimal points is fine.