Steps Toward Quantifying Galaxy Structure

Kory Kreimeyer
Mentor: Steve Odewahn
McDonald Observatory REU Program
Summer 2009

During the course of this summer at McDonald Observatory, I performed many various tasks related to the study of galaxy structure and morphology. Unfortunately, I was not able to attain any worthwhile results, but I now understand far more about the process and the steps involved. The steps that I performed this summer include: initial data reduction, the calculation of zero points, sky background subtraction, galaxy locating and sizing, and an attempt at fitting surface brightness profiles.

All of the data I manipulated this summer were taken either just prior to the REU program by my mentor or during the first few weeks of the program by the REU students. All images came from the 0.8 meter telescope at McDonald Observatory. It uses a prime focus corrector instrument (PFC) to gather data on a charge-coupled device (CCD) chip, which is the standard digital photography device for all modern telescopes. Also, three different filters were used to image each galaxy, providing data on their appearances in red, green, and blue light.

The first step in any astronomical image analysis is data reduction and calibration. During any observational run (of which the REU students performed several), a number of calibration frames are taken to improve the accuracy of the data. The first that I applied to the raw data were the "zero" or "bias" frames. These are calibration frames in which no exposure is taken. The data are simply read out from the CCD into an image, giving a map of the "read-out noise." Multiple zero frames were combined, and the large scale structure in this noise was removed from the science images.

I also applied a flat field calibration to all of the images. The flat field frames for this were also gathered during the observing runs by pointing the telescope at an evenly illuminated white sheet within the dome. Any variations in these images are then the result of varying sensitivities across different areas of the CCD chip and will be present in the science frames as well. I once again combined several of these calibration frames to find all these consistent variations and then removed them from the science data.

Having finished the main calibration of the images, I moved on to a simple method of calculating photometric zero points. A zero point is a parameter specific to each individual image and depends on the telescope and instrument used, the weather conditions and seeing effects at the time, as well as many other factors. Since all of these data were taken on the same instrument, however, I wanted to test how similar their zero points might be. This required knowledge of the actual magnitudes of the galaxies in the images, but I was able to find published values for almost all of these using internet database tools. After calculating all of the zero points, I averaged them for each filter and found very little deviation. In fact, these average values were used as part of later analyses.

My next undertaking was to thoroughly test a sky background subtraction algorithm. Accurately accounting for the sky signal in the images is one of the most important steps when attempting to study galaxy morphology, as the extended reaches of a galaxy can easily be twenty to thirty times dimmer than the night sky. To approach this, I first manually selected the areas of each image to be averaged, avoiding all stars and other objects. Then, I ran the program in the fully automated way and compared the two values for each image. I found that there was very strong agreement between the two methods. This verified the validity of the simple and automatic process for sky subtraction.

I next worked on a script to both locate a galaxy in an image and determine its size. I found that picking the brightest pixels of the image often detected stars, so my approach to locating the galaxy involved having the program search for a large area of only moderately bright pixels. This was still sometimes a troublesome prospect until I patched out all other light sources from the image by replacing them with local sky values. This turned out to be a trivial step to perform, but it improved the accuracy of the finding algorithm to essentially one hundred percent.

From the center of the galaxy, however, it was not a simple thing to measure size. My initial attempt was a flux-limited search outward from the center. This was often thwarted by gaps or holes in the structure of the galaxy and seemed to be a very ineffective way of gauging the size. Next, I attempted to look at the distribution shape of the flux across the galaxy. The main problem I encountered with this method was the extreme prominence of the bulge in relation to the disk. The width of the galaxy bulge was very easy to measure, but this value does not necessarily correlate with the visible size of the galaxy. Nevertheless, this was the technique I settled on using, with marginal success.

The final step in the analysis that I performed was to generate and examine elliptically averaged surface brightness profiles for galaxies. Using the galaxy center, size, and orientation data obtained from my program, I applied an elliptical annuli generating code to the images. From a single large ellipse outside the galaxy, it mapped a series of fifty concentric elliptical annuli of progressively smaller radii onto the galaxy. Within each annulus, is summed the flux and computed a generic magnitude. This information, magnitude versus radius, is a surface brightness profile.

Typically, these profiles are described by two components: an exponential fall-off for the disk of the galaxy, and a profile known as the $R^{1/4}$ law for the bulge. Structurally, the two functions are very similar, but the $R^{1/4}$ profile has a much sharper peak. I first attempted to fit these by converting the appropriate values to logarithmic scales to produce linear plots, but the values obtained from the two linear fits did not mesh well. I determined that it would be necessary to fit both components of the profile simultaneously and located a least-squares fitting algorithm that could theoretically perform this task. In practice, the fit was still quite difficult to create, as the program would preferentially choose only one component to fit the profile and would return unseemly values for the parameters of the other component. This inclined me to be less trusting of even the seemingly reasonable values it sometimes produced.

Unfortunately, I did not have time to work out the nuances of this problem; however, the final goal was definitely within sight. A comparison of the two components of the profile would give valuable information about the bulge to disk ratio, which is the primary feature used in the classification of galaxies. With a bit more time, this final stage could have been reached, but I also feel that a great deal of personal progress was made this summer anyway.