

Chapter 17

Finding the ages and distances of Milky Way globular clusters

As I described in Ch. 16, the size of the Milky Way can be probed using globular clusters. They lie above and below the Milky Way's disk and therefore are less affected by the thick dust that is concentrated there. In this project, we'll measure their distances to plot the Milky Way's size, and we'll measure their ages, to estimate an upper limit to the age of the Milky Way as a whole.

17.1 Data Sources and Software

If you want to begin with the images themselves, you can download Hubble images from the MAST Archive or Sloan images from the SDSS image server. For those who want to jump ahead to the catalog data, you can use the Hubble Source Catalog or the SDSS database, respectively.

The software needed for this section will be

- a FITS image viewer, such as DS9 (for those working with the images),
- an aperture photometry program, such as Aperture Photometry Tool (also for those working with the images),
- and a spreadsheet, such as Excel.

For readers who want to begin with reduced photometry data, skip ahead to section 17.3.

17.1.1 Obtaining Sample Images

For your first color-magnitude diagram, try the SDSS images for the globular cluster NGC 5466 that I have posted on the course's Blackboard site. Go to **Lab Instructions** → **Globular Clusters** and download the two files, `frame-g-004623-1-0324.fits` and `frame-r-004623-1-0324.fits`.

17.1.2 Obtaining SDSS Images

Below is a list of several globular clusters that are easily visible in the SDSS archive. Download the g and r filter images for each of these, as I describe in the following section.

	Name	RA (deg)	Dec (deg)
1.	M 2	53.370738	-35.769712
2.	M 3	42.207851	78.708428
3.	M 5	3.862575	46.795729
4.	M 13	59.009386	40.911864
5.	M 15	65.012586	-27.312821
6.	M 53	332.964972	79.765037
7.	M 92	68.338995	34.859424
8.	Pal 5	0.852	45.86004
9.	NGC 5466	42.15026	73.592027
10.	NGC 4147	252.848443	77.188751
11.	NGC 5053	335.6901	78.944163

Go to the SDSS Science Archive Server page to download centered images of your clusters.

- Start on the main page, `sdss.org`, and from the menu options, select
Data →
Imaging Data →
Science Archive Server/Explore SAS Imaging →
SDSS III Imaging: Single Field Search →
Search by Object Name
- Enter the cluster name and hit “submit.”
- Under the color preview image, select the *g*-band and *r*-band FITS images to download. We’ll use just these two filters in our analysis.
 - Note: If you’re on Windows, you may need to right-click and save to avoid having it open inside the browser.
- Save the files to your working directory.
- The files will arrive compressed in the BZ2 format and will need to be decompressed. If you have a decompression program installed (standard on Macs), you can simply double-click on them, and the FITS file will appear. If this doesn’t work, then you’ll need to get one of the many free programs that can handle BZ2. WINZIP is a popular one, and you can find others on the web.




Figure 17.1: The view of NGC 5466, as seen from the Sloan Digital Sky Survey in the green filter

Before going on, you should take a look at the cluster with the DS9 program (see Fig. 17.1), so that you can be sure you got the right image, and so you know what the cluster looks like and have an idea how good the image is.

17.2 Photometry

17.2.1 Introduction

Photometry is simply the measurement of an object's brightness in a given filter. What we measure is not the star's luminosity, which is an intrin-

sic physical property, expressed as either absolute magnitude or luminosity (usually in units of erg/s). Instead what we have is measurement-dependent, expressed in apparent magnitudes or units of flux, such as erg/cm²/s or W/m². For optical (visible-light) work, fluxes are usually given in erg/cm²/s but converted into apparent magnitudes, which are easier to work with.

The two main ways of measuring a star's flux are "PSF-fitting" and "aperture photometry," which we will do here. PSF-fitting means fitting each star we study with a model of the telescope camera's Point Spread Function, its diffraction pattern; that is, the blur pattern of any point-like source of light, like a star. This is more accurate than aperture photometry, but it takes a great deal of work, and I think we can get by with the results from the much easier aperture photometry method. Aperture photometry means measuring the light of each star through an "aperture," literally an "opening." In practice, we take the image, draw a circle around the star, and assume that the light within the circle comes from just that one star and from the background glow of the sky (see Fig. 17.2). Then we measure the brightness of the sky, subtract that much light from the aperture, and what's left should be from our star.

So how do we measure the sky brightness? We'll deal with flux per area, which in our case will mean flux per pixel. We need to measure the brightness in lots of pixels and average them together. In a properly-calibrated image, such as what you'll be using, the sky is pretty much "flat," meaning of even brightness, all across the image.¹ So we can measure the sky anywhere we don't have a star or galaxy showing, although we need to keep in mind that there could be some faint, extended light that is hidden down near the noise level and not obvious to the eye. We'll measure it within an annulus (a ring) centered on the star and a bit outside our aperture. We want the annulus to be far enough away from the star that we're not picking up too much light from the PSF "wings" (extended, faint light of the star's diffraction pattern), but not so far away that we're running into other stars. The wider we make the annulus, the more pixels are included, and the better the statistics are for measuring our sky. But if we make it too wide, we might be including other stars or regions that have somewhat different sky levels, since the flatfielding isn't always perfect.

¹The raw images from a CCD camera show a lot of variation in their sensitivity from one side to another, but the observatory corrects for this by "flatfielding" the image, making the sky, well, flat. The images you'll be using have all been flatfielded.




Figure 17.2: Aperture photometry. The red circle encloses the aperture for measuring the star brightness. The sky brightness is measured in the annulus (ring) between the green and yellow circles, and that level is subtracted from the light in the aperture.

17.2.2 Making Source Lists and Calculating Photometry

To identify the stars in the cluster and make a list of them, we use the Aperture Photometry Software program. First, we'll load the g -band image and set the photometry for this filter.

- Within the APT program, open the g -band image for the cluster.
- Go to `More Settings`, which opens up a window for the photometry settings.
 - In this window, set the sky algorithm to Model B (Sky-annulus median subtraction) or Model F (Non-sky-annulus local mode subtraction).²

²Model B is the standard way of subtracting the sky that I've described above. It works well when most of the field shows the sky, and stars have gaps between them. But in crowded fields, it will be hard to find the sky level for all of the stars in the way,

- If this is an SDSS image, set zeropoint to 22.5. For other observatories, see the table of zeropoints.
- Hit `Apply settings`.
 - * If you chose Model F, it will ask if you want to compute the sky background now.
 - * Select `Yes`, and it will bring up the `Local Non-Annulus Sky-Background Tool` window. There are options for adjusting the statistics, which we won't fool with now.
 - * You should select `Compute`, and then close the `Sky Background Tool` window.
- Close the `More Settings` window.
- On the main window, click on the “Big Button on Right.” This opens the `Elliptical Aperture Attributes` and `Photometry Major Radii` windows.
 - Make sure the checkbox is selected for a circular aperture (this is the default), and set “Major radius” to 10 pixels. This sets the radius of the aperture.³
 - Hit `Apply` and close the `Elliptical Aperture Attributes` window
 - In the `Photometry Major Radii` window, you can change the size of the sky annulus. I usually find the default settings pretty good. If you change them, select `Apply` before you close the window.
- Select `Recompute Photometry`, which applies everything you've set so far and tells APT how to compute the star magnitudes.

Next, we're going to make our list of sources (stars) and compute their magnitudes.

and Model F is better. This scans the entire image and looks for bits of sky background between the stars, making a map of the sky level across the field.

³It is possible to make an elliptical aperture by unselecting the circular aperture checkbox and entering both a major and minor radius. This sets the semimajor (longer) and semiminor (shorter) axes of the ellipse. An ellipse also needs a rotation angle specified. The angle is that of the ellipse's major axis measured counterclockwise from the $+x$ axis. We will not use elliptical apertures for stars, since they are round, but we apply it to galaxies, as shown in later chapters.

- Select `List Results` to see if a list of photometric results is already in memory. If any list shows up, clear it and then close the window. Be sure to clear the results before each run, or they'll get appended to the bottom of the list and be unusable.
- Select `Source List`. This opens the `Source-List Tool` window, which allows you to select or create a list of sources in the image.
 - If you already have a list of sources for this cluster (perhaps from an image in a different filter), then skip the indented instructions below. If you don't have a list for this cluster, you'll need to create one:
 - * Select `Pixel` coordinates. This will make the list be the (x, y) coordinates of the stars. In some cases, we'll make a version in equatorial (RA, Dec) coordinates, but that comes later.
 - * Select `Create Source List`, which brings up the `Source-List-Creation Tool` window.
 - * Select `Reset Source-List Filename`, which lets you choose the name of the source list and where it will be saved. I recommend giving it a name that makes it easy for you to figure out. For example, if this list is made for the g -filter image of the cluster NGC 5466, using pixel (x, y) coordinates, then name it `ngc5466-g-sources-xy.txt`.
 - * The other options on the page deal with the statistics on how bright above the noise level the pixels should show to be counted as stars. You can leave these at the default settings.
 - If you already have a list of sources for this cluster, select `Choose Source List` to load it. Select either `Pixel` or `Equatorial` coordinates, depending on the format used in your list (see the discussion at the end of this section for which one is appropriate).
 - Select `Create Source List`, which will make APT scan through your image and make the list. This will take a few seconds. After it says `Processing completed`, select `Close window`.
 - In the `Source List Tool` window, which will still be open, make sure it shows the name of the source-list file that you just made. It should show the number of sources in the list. If it reports "0,"

there's a problem, and you should adjust the statistical settings in the `Source-List-Creation Tool` window.

- If there are sources detected, then select `Automatically process source list` and let it run. Depending on the number of stars in the list and the power of your computer, it could take anywhere from a minute to a few hours. On my 2015-era MacBook Pro, a list of 2,388 stars takes about one minute. If you have selected the `Check box to overlay sources on image`, the image display in the main window will show green circles around each of the stars it has identified (See Fig. 17.3). This is a good way to confirm that things are working properly.
 - Once it's finished, select `Close window`.
- Now select `List results`.
 - The results show a lot of information. Each row is a single source—one star. Each column lists a different piece of information about that star. You have the coordinates (both x, y and RA, Dec) for the center of the aperture and for the “centroid” (a kind of weighted average for the center of the star itself). There's the source intensity and its uncertainty, given in the units used by the image. For SDSS images, the unit is the adorably-named “nanomaggies.” Other observatories might use “counts” (number of electrons counted for the entire exposure) or electrons per second (e-/s). Then there is the star magnitude and its uncertainty. There are lots of other details, but all that concerns us is position and brightness.
 - Check that there is the same number of rows as you had sources in the list. If you have more rows than sources, you may have had an earlier set of results still in memory that you forgot to clear. Clear the results and process the list again.
 - If all is well, you'll save the list in the “Comma-Separated Values” (CSV) format, which can be read by a spreadsheet. Select `Export Contents`.
 - * Select `Reset CSV filename`.

- * Choose a directory and filename. For the filename, use the cluster name, followed by the name of the filter, followed by “phot” for photometry, and use “.csv” for the file suffix. So our current example would be titled, `ngc5466-g-phot.csv`.
- * Select .
- Select to erase the results from APT. Don’t worry—they’re saved in the CSV file you just created.

That’s it for this filter. For making a color-magnitude diagram, we must measure photometry on the cluster in *two* different filters. The red (*r*-band) filter is the next one used in my example, but the process is the same for any other filter.

Now we’re going to set things up to run the next image. The first task is to create a source list for the *r* filter. The stars are in the same place in space no matter what, but if our two images are not perfectly aligned, they won’t appear in the same place in the two pictures. That is, the stars’ (x, y) coordinates might not be the same. You can check how well aligned they are by opening both the *g* and *r*-band images in two separate frames in DS9 and letting them blink. If they’re off by a significant fraction the width of your aperture, then you must make a source list using the equatorial coordinates, RA and Dec. If they’re well aligned, then you’re fine to use the existing pixel-coordinate list from the *g*-band image.

To make an equatorial-coordinate list⁴, open the photometry table you saved from the *g*-band image in Excel. Copy the columns `ApertureRA` and `ApertureDec` into a blank text file, and save it as `ngc5466-g-sources-equ.txt`, where “equ” tells you it’s in equatorial coordinates. You should have two columns, separated by tabs or spaces, but not by commas. The column headings are fine to leave in. If you have a sophisticated text editor, like Emacs, you can copy and past these columns directly, without having to go through Excel, but you’ll need to replace all of the commas with spaces.

Now run APT on the *r* filter image. When you load the new image, APT will ask you what you want to do with the current image. You should

⁴Your image must have a World Coordinate System (WCS) set up in its header. If not, the website astrometry.net can be used to put a WCS into the header. You’ll need to create a WCS for both of your images and re-run the photometry, so that the table will have RA and Dec.

clear the current image (option A). Now repeat the steps in §17.2.2⁵, but use the equatorial-coordinate *g*-band source list, rather than creating a new one. This will ensure that you measure the same stars in the same order in both filters. If you didn't, then we would not easily be able to make the color-magnitude diagram for the cluster in the next section.

17.3 Color-Magnitude Diagrams

17.3.1 Introduction

17.3.2 Constructing the Diagram

Now that you have lists of the cluster's star magnitudes in two filters, we'll use those to make the color-magnitude diagram. Open your Excel or another spreadsheet program and create a new, blank form. Open the *g*-band photometry table (or whichever filter is the shorter wavelength of your pair) in Excel and copy the **Number** and **Magnitude** columns into the first two columns of your new form. Retitle **Number** as **g Number** and **Magnitude** as **g**, because it is standard (and convenient) to represent apparent magnitudes with the filter name. Now do the same for the *r*-band table, putting them in the third and fourth columns. Retitle its number as **r Number** and its magnitude as **r**.

Look at the source numbers for the *g* and *r* tables. Ideally, they will match up, line for line, all the way to the bottom. But if the *r* image is shifted with respect to the *g* image, then some of the sources found in *g* may not show up in *r*, creating a mismatch. Everywhere you see numbers skipped in the *r* list, select the **r Number** and **r** cells for the next source and have Excel **Insert Cells**, shifting the selected cells down (Fig. 17.4).

We should not leave the magnitudes blank, even for the rows with missing stars, but Excel has a handy function for this that identifies something as "Not A Number." You may see this as "NaN" in some programs. In Excel, go into each blank cell in the magnitude column and enter =NA(). Once

⁵**Note:** In earlier versions of APT, I have seen the list processing hang for hours without ever completing. I believe this was caused by it using up too much memory while processing the second filter. If that happens to you, then you should quit APT after the first filter and restart the program before running the second filter. I have not seen this bug in newer versions.