MIT Invitational, Jan 2019

Astronomy C: Free Response Walkthrough

Competitors: ______________________________

School name: ______________________________

Team number: ______________________________

INSTRUCTIONS

1. Please turn in **all materials** at the end of the event.
2. You may separate the pages, but do not forget to put your **team number** at the top of all answer pages.
3. Write **all answers on the answer pages**. Any marks elsewhere will not be scored.
4. Do not worry about significant figures. Use 3 or more in your answers, regardless of how many are in the question.
5. Please **do not access the internet** during the event. If you do so, your team will be disqualified.
6. Good luck! And may the stars be with you!

Section B

Each subpart is worth 2 points. Be sure to use image sheet B. For calculations, just write the final answer (be sure to answer in the units requested).

1. Image 1 shows a JS9 visualization of data collected by the Hubble Space Telescope.

   (a) What is pictured?

   The antennae galaxies. This picture is probably not easily found on Google, so it requires you to be able to know what the general structure of this antennae galaxies looks like, and recognize that in this image.

   (b) Image 2 shows the x-projection (taken via arithmetic mean) of the data in the green box. The x-axis of the graph represents the pixel number, starting from the left edge of the box. What does the y-axis measure?

   The average number of photon counts. Telescopes use digital detectors which count individual photons. Telescope data basically consists of information about how many photons were counted in each pixel of the image. The number of photons corresponds to the how bright the light is, so I also accepted "luminosity" as an answer. For more intuition on how this works, play around with the JS9 tool at `https://js9.si.edu/`. As you mouse over the image, there should be a green number in the top-left corner; this is the photon count at the pixel you're hovering over!

   (c) Observe the color scale at the bottom of image 1. What kind of scaling function is being used (i.e. linear, quadratic, cubic, logarithmic, sinh, or something else)? Why might this kind of scale make sense for astronomical data?

   Logarithmic. Astronomical energies span many orders of magnitudes; a logarithmic scale will get a sampling of all these different orders of magnitude. You've seen many examples of log scaling: for example, the magnitude scale is based on a logarithm of luminosity. The HR diagram very frequently scales the temperature axis logarithmically, as well as the luminosity axis. For more information, see here.

   A common answer was exponential, because the numbers on the scale grow exponentially. Keep in mind that the numbers on color scale (i.e. the photon counts) represent the *input* to the color function. If you imagine plotting points at all those x-values, at uniformly increasing y-values, you'll get a logarithmic function. Another perspective: since input is growing exponentially related to the output (i.e. $x = 2^y$), you have to take the inverse, which gives you a logarithm. These are just a bunch of different ways to reason through the problem if you didn't already know that logarithmic scales are very common in astronomy (and science in general).

   Keep in mind, of course, that this is a false-color image. In other words, this whole color scale business has nothing to do with the ACTUAL color of the light. It's just a way for us to make the image easy to look at, and easily find areas of high brightness (i.e. high photon count). This leads in nicely to the next question...

   (d) What band of light was most likely captured in this image? Hint: consider the telescope.

   Visible (aka optical). The Hubble telescope can take images in visible light, near infrared, and ultraviolet, but it specializes in taking visible-light images. Even if you didn't know that, you might be familiar with various images of the antennae galaxies, and know that the Hubble images were taken in the visible waveband.

   Even if you still didn't know that, you can kind of make an educated guess based on the image; it's bright at the core and in the arms (suggesting UV/visible/infrared), but is dark in the star-forming regions (eliminating the infrared option - this is related to part e). The fact that the core is much brighter than the arms makes it unlikely that it was UV; the cores of galaxies tend have mostly older stars, and older stars are not bright in ultraviolet. (Since hot blue stars die young, older stars must be redder, and therefore faint in UV light).

   In general, understanding multiwavelength astronomy is super super important. Here's 4 different resources that may help: basic, basic, intermediate, and advanced.

(e) After messing with the color scale some more, you end up with image 3. Your friend notices the dark region encircled in image 3 and says, "If you had taken this image in infrared, that would be one of the brightest parts in the image!" Is she right? Why or why not?

Yes. The dark region is a large swath of dust and cold gas which is obscuring a site of intense starburst. It absorbs the new stars' high energy light (including visible wavelengths) and re-emits the absorbed energy as infrared. So it would appear dark in visible light and bright in infrared.

In terms of observation, dust is the enemy. It absorbs all sorts of wavelengths of light, especially the interesting high-energy wavelengths. If you ever see an unusual dark region in an image, your first suspicion should be dust - large amounts of it, blocking out the light behind it. In this case, we pair this knowledge with the fact that dense clouds of gas/dust are ideal spots for star formation, and the fact that the antennae galaxies are starburst galaxies. This strongly suggests that there is star formation going on deep inside the dust cloud, producing hot, bright stars.

Even if there wasn't star formation going on inside the cloud, dust clouds are generally infrared-bright because they re-emit any light that they absorb in infrared. To be more precise, dust grains absorb high-energy light, which heats up the grains. The slightly-hotter grains now emit blackbody radiation of their own. But they're still nowhere near as hot as the stars that originally emitted the light. As a result, they emit lower-energy radiation: infrared.

For more information about dust and infrared astronomy, here are some links: basic, intermediate, advanced.

2. Consider a binary system involving an evolved star and a black hole. The average separation between the two is determined to be 3 AU, and the orbital period is 0.8 years.

(a) What is the total mass of the system, in solar masses?

42.2 solar masses.

Using the general version Kepler's 3rd law, in solar units:

$$M = \frac{a^3}{P^2} = \frac{27}{.64} \approx 42.2$$

Remember that you don't need all the constants ($G/4\pi^2$) when you're using solar units; those only come up when you're using the MKS unit system (meter/kilogram/second).

(b) How hot (in Kelvin) does the accretion disk have to be if the thermal radiation peaks in the X-ray band? (A range of answers will be accepted.)

2.9×10^5 to $2.9 \times 10^8 K$

We generally take X-rays to have wavelengths between 0.01 to 10 nanometers. Taking a representative wavelength $\lambda = 1$ nm, we can use Wien's displacement law:

$$T = \frac{b}{\lambda} \approx \frac{2.89 \times 10^{-3}}{10^{-9}} = 2.89 \times 10^6$$

(c) The temperature of the disk depends on distance from the black hole: $T \propto r^{-1/2}$. Your friend predicts that the majority of the disk luminosity will come from the outer regions of the disk, since there's more gas to emit radiation. Is she right? Explain why or why not.

No. Although there's more gas, the inner part is hotter, and since the luminosity is dependent on T^4 (Stefan-Boltzmann law), the temperature overcompensates for the lack of radiating material. The inner part is in fact more luminous.

If you're super advanced, you can use elementary differential calculus to show that the luminosity actually decreases with radius (one team actually did this! Good job!). Of course, we don't expect you to know calculus. Many teams tried to do naive analysis and ended up concluding that the luminosity is constant throughout the disk. This isn't right, but is an understandable error, so I gave credit if you had the right idea.

If you're interested, look here for a pretty advanced resource on accretion disks.

3. Determining the distances to galaxies is important is astronomy, and there are a lot of different techniques to do it.

(a) Consider a Type I Cepheid with a period of 25 days, and an average apparent magnitude of 19.1. Using image 4, estimate its distance, in parsecs.

780,000 pc ± 100,000 pc

Using the relation in image 4, the luminosity is about 11,000 solar luminosities. We want to use the distance modulus to estimate the distance, so we need the star's brightness in terms of absolute magnitude, not solar luminosities. We can use the relation between bolometric magnitude and luminosity:

$$M_{bol} = M_{bol,\odot} - 2.5\log_{10}\left(\frac{L}{L_\odot}\right) = 4.74 - 2.5\log_{10}(11000) \approx -5.36$$

Then, the distance modulus:

$$d = 10 * 10^{(m-M)/5} = 10 * 10^{(19.1-(-5.36))/5} \approx 780,000 \text{ pc.}$$

(b) Suppose the measurement of the peak apparent brightness of a type Ia supernova had an uncertainty of ±0.5 magnitudes. What is the corresponding uncertainty in distance? Express your answer as a ratio between the upper and lower error bounds.

1.58 (reciprocal 0.63 for half credit)

Here, we want to take the ratio between two distances, so we need to do distance modulus twice and take a ratio. As it turns out, this ratio doesn't depend on how far the supernova actually is; in other words, you can get the answer by plugging in dummy numbers and getting the numerical result. Here, for the sake of completeness, I'll derive the result symbolically.

We'll call the lower error bound d_1 and the upper bound d_2. Then

$$d_1 = 10 * 10^{(m-M-0.5)/5} \qquad \text{and} \qquad d_2 = 10 * 10^{(m-M+0.5)/5}$$

where m is the measured magnitude, and M is the absolute magnitude. Then we take the ratio and use exponent rules to simplify:

$$\frac{d_2}{d_1} = \frac{10 * 10^{(m-M)/5} * 10^{0.5/5}}{10 * 10^{(m-M)/5} * 10^{-0.5/5}}$$

Note that the first two factors (including the one with our magnitudes in it) cancel out! This leaves us with a numerical result:

$$\frac{d_2}{d_1} = \frac{10^{0.1}}{10^{-0.1}} \approx \frac{1.26}{.794} = 1.58$$

(c) Imagine that you tracked a type Ia supernova and determined its distance. But then, your friend (who works at a gravitational wave lab) says that he measured a gravitational wave "chirp" originating from the supernova. (Let's say that this is in the future, where gravitational wave detectors are much more sensitive than they are today.) Is this information relevant (i.e. should you modify your distance calculation)? Why or why not?

Yes, the original estimate was probably an underestimate. This was likely a double-degenerate scenario.

Many teams considered how gravitational waves might affect the light produced by the supernova, and concluded that they don't (which is true). This is a good first step, but you need to also consider what gravitational waves say *about* the supernova. Gravitational waves are produced primarily by two dense objects spiraling around each other at high angular frequencies; a "chirp" implies that the two objects actually merged (listen here).

Mergers in the context of Type Ia supernovae should remind you of the double-degenerate model. In this scenario, the white dwarf collides with another dense object (as opposed to accreting matter from a star). In this case, the white dwarf undergoes the same explosion (so it's still a type Ia supernova), but the total mass might be much more than the Chandrasekhar mass (since there was no steady accretion process, it was a sudden mass dump). This violates the central assumption of using Type Ia supernovae as standard candles! If the mass was not the Chandrasekhar mass, then the absolute luminosity is not known, and we can't use the distance modulus.

This question was hard because it required you to synthesize different information that isn't explicitly mentioned in the problem. Very few teams got this question right, so don't feel bad if you didn't!

(I should mention: this question is going off of the traditional theory for Type Ia supernovae. Over the past couple of years, research has suggested that the story isn't quite as straightforward as it once seemed. In my opinion, though, the cutting edge research about Type Ia supernovae is too uncertain, and is out of the scope of the Astronomy event.)

(d) The 21-cm line of a distant spiral galaxy is shown in image 5. (Warning: read the x-axis carefully.) Estimate the distance to the galaxy, in megaparsecs. Hints: The 21-cm line has a wavelength $\lambda = 21.106$ cm. Use $H_0 = 72\ \mathrm{km\,s^{-1}Mpc^{-1}}$. Be sure not to round too soon!

12.4 Mpc $\pm$ 1 Mpc

Notice that the spectrum has a doubly-peaked line. This is evidence of disk rotation; the part of the galaxy which is rotating towards us is blueshifted (relative to the galaxy), while the part rotating away is redshifted. The "actual" line would be right in between the peaks, at 1416.2 MHz, or $\lambda = 21.169$ cm (using the wavelength-frequency relation, $\lambda\nu = c$). Notice that this wavelength is longer than the normal 21-cm line – this tells you that the galaxy is redshifted, and you can use Hubble's law. The recessional velocity is given by the Doppler equation

$$v = c * \frac{\Delta\lambda}{\lambda_0} = c * \frac{.063}{21.106} \approx 894 \text{ km/s}.$$

Then we use Hubble's law:

$$d = \frac{v}{H_0} \approx \frac{894}{72} = 12.4 \text{ Mpc}.$$

This question was strongly inspired by this resource.

Many teams got answers that were close but not close enough. I have a suspicion that this comes from rounding the speed of light to 3×10^8 m/s. Usually this is fine, but this question required numerical precision.

(e) Using both image 5 and the Tully-Fisher relation shown in image 6, estimate the mass of the galaxy, in solar masses.

Approximately 10^9

We want to use the Doppler effect again, but this time not the recessional Doppler, but the rotational Doppler. This will tell us how fast the galaxy is rotating. Using the frequency formulation of the Doppler effect, we get

$$v_r = c\left(1 - \frac{\nu}{\nu_0}\right) \approx c\left(1 - \frac{1416.6}{1416.2}\right) = -84.7 \text{ km/s}$$

On the graph (which is logarithmically scaled!), this corresponds to a mass of about 10^9.

Section C

Each subpart is worth 3 points, unless otherwise specified. Be sure to use image sheet C.

4. A luminosity function is a "distribution" representing how many objects exist at a certain luminosity, with the amount on the y-axis and the luminosity on the x-axis. Image 1 shows a simulation of luminosity functions of white dwarfs for a galaxy at different ages between 8 Gyr and 16 Gyr, each depicted as a different curve. Notice how the curves show practically no dependence on age towards the beginning but diverge considerably at lower luminosities.

   (a) As white dwarfs age, what happens to their temperature and luminosity?

   Both the temperature and luminosity decrease.

   (b) Which curve (A or B) represents the galaxy when it is younger?

   Curve A

   (c) Curve A has a distinct lack of extremely dim white dwarfs, while the Curve B has significantly more. From a stellar evolution perspective, briefly explain why.

   It takes time for white dwarfs to cool and become that dim. At the age represented by Curve A, the white dwarfs in this galaxy haven't cooled as much, but at the age represented by Curve B, they've had the time to cool much more. Since the luminosity of an object is proportional to T^4, a lower temperature would mean a dimmer white dwarf.

   (d) In the interval $0 > \log(L/L_\odot) > -4$, the luminosity function appears to be strictly increasing. Why is this so?

   The cooling rate slows as the white dwarf cools (the graph of temperature vs. time would appear to be concave up and decreasing), until crystallization sets in. This causes the amount of white dwarfs at that luminosity to "build up"

5. The Schmidt law is an empirical relation between gas surface density and star formation rate (SFR) in galaxies, first examined in 1959 by Maarten Schmidt. In his paper, he suggested that gas density and SFR are related as such:

$$\Sigma_{SFR} = (\Sigma_{gas})^n$$

   Data from 15 galaxies (some normal spirals and some starbursts) have been given in image 2.

   (a) Based on the data, would you expect n to be greater than 1 or less than 1?

   Greater than 1

   (b) Give a brief qualitative explanation for the physical basis of Schmidt law. Based on this, would you expect starburst galaxies to have higher or lower gas surface densities than their normal spiral counterparts?

   Higher surface gas density results in a higher chance for higher mass clumps that can exceed the Jeans mass, which would lead to runaway contraction and heating, marking the beginning of star formation. Based on this, we would expect starburst galaxies to have higher gas surface densities.

   (c) The SFR for the starburst galaxies in this sample were derived from measurements of their FIR luminosities (LFIR). Explain why astronomers can use FIR luminosities to derive SFRs for starburst galaxies with reasonable accuracy, but not with normal spiral galaxies.

   In normal disk galaxies, the relationship between the FIR luminosity and the SFR is complex because stars with a variety of ages can contribute to the dust heating, and only a fraction of the bolometric luminosity of the young stellar population is absorbed by dust. However, in starburst galaxies, the physical coupling between the SFR and the IR luminosity is much more direct. Young stars dominate the radiation field that heats the dust, and the dust optical depths are so large that almost all of the bolometric luminosity of the starburst is reradiated in the infrared.

(d) (6 points) Image 3 shows a comparison of FIR and Brγ-derived SFRs for starburst galaxies. The solid line shows the correlation expected if the two sets of SFRs were equivalent. An astronomer looks at this plot and notices that the FIR-derived SFRs appear to be higher than their Brγ-derived counterparts.

i. What is a physical limitation of Brγ which could account for the discrepancies between FIR-derived SFRs and Brγ SFRs?

> The extinction in most star forming regions is so large that one expects part of the ionizing radiation from the starburst to be absorbed by grains, and in some objects, extinction of Brγ itself is probably significant, making it seem dimmer, and as a result, less active (star formation wise) than it actually is.

ii. What is an observational limitation which could account for the discrepancies between FIR-derived SFRs and Brγ SFRs?

> If the starbursts are observed after the peak of the burst, then the dust heating is dominated by longer lived stars and the effect of the Brγ emission line is less than it actually was.

6. One way of helping answer the age-old question of "where did we come from?" is to observe the progenitors of our own galaxy, the Milky Way. In order to look far back in time, we have to look for galaxies far away, which can get difficult. One creative way of finding these distant galaxies is by looking for objects that seem to "drop out" when we image them in different filters (wavelengths). Image 4 shows the spectra of two distant galaxies, labelled A and B.

(a) Both galaxies are receding at very high velocities due to the expansion of the universe, which causes their spectra to get redshifted. If their radial velocities relative to us were zero, at what wavelength would we notice the "break" in the spectra? What causes this "break" in the spectra?

> We would see the break at 91.2 nm. Any photon emitted with wavelength shorter than the Lyman Limit (91.2 nm) be completely absorbed by hydrogen gas both in a galaxy and along the line of sight to us instead of reaching our telescopes, resulting in a "break".

(b) Based on image 4, what are the redshifts of each galaxy? Which one is further away?

> Galaxy A: 3.1-3.5; Galaxy B: 4.1-4.5. Galaxy B is further away.

(c) An astronomy student attempts to use this method to investigate nearby galaxies with a ground-based telescope. They search for a "break" in the spectrum, like the one shown in image 4, but are unable to find any useful information. Why is this so?

> The galaxy's spectrum would have a break in the UV region at 91.2 nanometers, as specified in (a). However, the Earth's atmosphere is extremely opaque at these wavelengths, so we would not be able to see anything. The reason we can see the breaks for these distant galaxies is because they have been redshifted to wavelengths that we can see through the Earth's atmosphere. Since these galaxies in this part are nearby, this cannot be the case.

(d) Let's put it all together. Suppose you examine a patch of the sky and collect the pictures shown in image 5, each taken through different photometric filters, denoted by the letters U, B, V, and I. Estimate the upper and lower bounds for the distance to this galaxy (shown by the yellow arrow), in megaparsecs.

> 3650-4100 Mpc