A Personal Approach to Research in Astrophysics

Cole Miller

The step from successful coursework to successful research is often a large one, particularly when the research is in astrophysics. There are a number of approaches I've picked up over the years that seem helpful to me. The disclaimer is that there are many different ways to do research, and many people may disagree with me; also, I am speaking from my experience in theoretical astrophysics, and other techniques may be useful in observational astronomy.

1. Talk to people

When one is doing a homework assignment, the point is often to do it alone. When doing research, the point is to solve the problem. If you spend weeks sitting in a corner banging your head against a wall because you aren't making progress, you aren't using your time efficiently.

Q: I want to talk to my advisor about my work, but he/she is unapproachable/busy/absent. What then?

A: Talk to postdocs or older grad students. When I started doing research as a grad student, I was embarrassed to talk to my advisor because he was able to answer all my questions without appearing to think about it; I felt I was asking low-level questions and making a bad impression. I resolved this finally by deciding that when I had a question, I would ask some postdocs and if *they* didn't know the answer, then I would ask my advisor. This way, the questions that got to him went through a filter.

Even if the person you're talking to doesn't help you directly, the process of explaining your problem is often a great help. I can't stress enough how bad an idea it is to curl up in an extradimensional ball and not interact!

This extends to conversations with people outside your university. Down the line, it will be important that professors at other places know who you are and what your work is. The best way for you to accomplish this is through personal contact. Your advisor should facilitate this by allowing you to go to conferences and make presentations and, if possible, by making personal introductions. You can also initiate this contact by calling or e-mailing experts in your field with various questions. Here you really don't want to get off to a bad start, so at least at first you should consult with your advisor about such external questions, then make contact (e.g., "Hello, Professor Jones, my name is Pat Johnson. I'm a graduate student working with Robin Smith, and I had some questions about axion stars. Do you have a few minutes to talk?"). If you make several such contacts with a professor, the professor will know who you are and will, hopefully, have gained respect for you.

2. Don't begin your research by reading the literature

There are at least two good reasons to not *start* a problem by reading the literature. First, you will understand the physics issues much better if you've wrangled with them yourself. Second, you have a worse chance of coming up with something really original if you contaminate yourself with the thoughts of others.

Of course, during your project there will come a time when it is very important to make a thorough survey of the literature. You can't think of everything by yourself, so looking at other people's work will give you lots of ideas. The other reason an eventual literature survey is important is that when you write a paper you don't want to tick people off because you haven't referenced their papers.

Nonetheless, the long-term benefit will be much greater if you spend a while (days, maybe) trying to work things out yourself from first principles. This is one of the best ways to develop intuition: apply it from the start and compare it with what other people have done.

3. Don't believe everything you read

When you do get around to reading the literature, don't treat it like holy writ. This is another big difference from taking classes. A textbook is often presented as the revealed truth, which you should learn from but not challenge. Papers are a different story. They are reports on current research and therefore haven't been tested fully and their arguments aren't as crisp as the ones given in textbooks. So, when reading a paper don't assume everything in it is right! One approach is to read just the abstract, then go off and think about things yourself for a while. Then, come back and read the introduction (for context) and conclusion (to see what the authors think they've added). If the paper is really interesting to you, read the body, but don't get bogged down in details. Above all, it is *essential* to read everything critically! As you read, constantly ask yourself: "Does this make sense?", "Does it conform to my previously developed intuition?", "If the authors disagree with me, what is the resolution?".

There are lots of bad papers out there. If you put your brain on autopilot and believe everything you read, you will mislead yourself.

4. Don't believe everything your computer spits out

A large fraction of modern astrophysics involves numerical computations. There are tons of different techniques: Monte Carlo codes, radiative transfer, hydrodynamics and MHD, and so on. The single cardinal rule that applies to all numerical work is GIGO: Garbage In, Garbage Out. The computer only does what you tell it to do, and it will happily spew out definitive-looking drek. The way around this comes back to intuition: what *should* the output be? If it doesn't conform to your expectations, find out why. It could be that your intuition wasn't developed enough, and that the output is right. However, in the majority of cases I've encountered, if the output isn't what you expect, it's wrong. I suggest that a productive attitude is that the computer is wrong until it is proven right! The point here is that although you use a computer to explore physics that you can't solve directly, you must run enough test cases that you can solve directly so that you trust the final output.

The tests you should conduct are both qualitative and quantitative. If you expect a x^2 dependence and the computer gives you a x^4 dependence, you need to know why. If you think the computer should give you 7.5000 and it gives you 7.48, you need to know why. Apparently minor inaccuracies for simple problems can magnify themselves in unpredictable ways in the more complicated problems you really want to solve.

5. Learn how the answer to a question depends on what you need

One of the toughest parts of the transition from classroom physics to research astrophysics is that the problems typically aren't solved in the same way. In a classroom, a physics problem is posed, then solved. In astrophysics research, the problems are often inherently far too complex for there to be an exact solution. Thus, approximations are necessary, and what approximations are justified often depends on what answer you need, to what level of accuracy.

Here's an example. Let's say I tell you that I have a filled water tank that has dimensions of 117 cm by 93 cm by 105 cm. If I ask you the mass of the water in the tank to the nearest gram, you have to not only multiply the dimensions together to get the volume, but also ask me the temperature of the water and the local atmospheric pressure, since the density changes with such quantities to more than a part in a million! If instead I ask you if you can lift the tank, you can immediately answer "no", because you can see at a glance that the mass is approximately a metric ton, which is beyond your capabilities unless you were born on Krypton.

Learning to see to the essence of a problem in this way is at the heart of much astrophysics research. Don't get bogged down in details unless you have to. Be lazy; use the simplest methods available, since they are also the most convincing!

6. Learn how to estimate

Even in the realm of arithmetic, simplicity is often a helpful guide. Estimates can come in different forms. Order of magnitude estimates are helpful when you want to know a quantity. They often require a fairly broad knowledge of helpful astronomical quantities, to be most effective. For example, let's say I ask you how far away a binary with a separation of 3×10^{13} cm could be resolved by the Hubble Space Telescope (with an assumed resolution capability of 0.02"). You could get out your calculator and multiply 3×10^{13} cm by the number of degrees in a radian and the number of arcseconds in a degree, and divide by 0.02. Or, you could note that 3×10^{13} cm is about 2 AU, and that a parsec is a parallax second with a baseline of 1 AU, to estimate that the distance is about 100 parsecs.

The other type of estimation is useful when you want to make an argument. In this type, you make generous assumptions and then show that whatever you want to prove or disprove still holds. For example, say that I tell you the estimated power of cosmic rays in our galaxy exceeds 10^{40} erg s^{-1}, and ask you if the ultimate source of the cosmic rays could be the magnetic fields of neutron stars. You could try to calculate the spectrum, which would be tough and unconvincing. Or, you could note that the birth rate of neutron stars is not greater than one per ten years, or 3×10^{-9} s^{-1}. Assuming very generously that the average surface magnetic field strength of neutron stars is 10^{15} G, the magnetic energy per neutron star is about 10^{18} cm^3 $\times$ 10^{29} erg cm^{-3} $=$ 10^{47} erg, giving a maximum magnetically powered energy generation rate of 3×10^{38}erg s^{-1}, far below the requirement. Since magnetic fields of neutron stars fall short even with the assumed high rate and magnetic fields, it ain't so.

As you can see, putting together such an argument requires that various numbers be at the tips of your fingers, such as the volume of neutron stars or their estimated magnetic fields in the example above. This leads to the next point.

7. Assemble a toolkit allowing for rapid assessment of ideas

Astrophysics involves a broader range of important physical concepts than any other field of physics. You can't expect yourself to derive everything from first principles every time you ask a question. It is therefore important to do many exploratory calculations to get an idea of what is important when. For example, which is a more important source of energy for a compact object over a long time, accretion and the release of gravitational energy or thermonuclear burning? It depends on the situation: for neutron stars accretion produces 20-30 times as much energy per gram as does fusing hydrogen all the way to iron. For white dwarfs, the ratio is reversed. What are the dimensions of the object of interest? What is the mass? What are typical luminosities, cross sections, time scales? Knowing the answers to such questions gives you a high-level feel for the problem that will allow you to evaluate new ideas much more rapidly than you could otherwise.

As before, a good way to make such ideas stick is to do the calculations yourself when possible, instead of just reading about them. You also have to know the regime in which the calculations are invalid. In the example above, if you are used to ignoring the release of gravitational energy because you deal with white dwarfs, you'll make errors when you carry your toolkit along to neutron stars!

Part of this toolkit is that you should have a thorough understanding of the simple ideas that are central to your field. For example, why are most X-ray bursts from neutron stars thought to be the result of thermonuclear flashes, even though accretion energy dominates over long times? Why are the microwave background photons scattered by gas in clusters thought to *gain* energy on average? These are the bedrock ideas on which you can build more speculative concepts.

8. Know the observations

Astrophysics is observationally driven. In almost all subfields of astrophysics, there are hordes of observational trends or individual sources that are not well explained by current theory. In order to explain these observations, you must know what they are! This is also a good way to make yourself popular with observers, who often rightly complain that theorists ignore much of the observational literature and results, and treat observers like unskilled labor. Observers and theorists need to work together to produce the best results, so from a theorist's perspective it is important to (1) know what observations have been reported, and (2) know how the observers arrived at their results. Point number (2) is that cutting-edge observations are often done at the hairy edge of detector and statistical capabilities, so don't assume that all reported observations are immutable fact! If you do, you'll end up chasing statistical or instrumental chimeras. If you ignore observations entirely, you'll end up traipsing through a theoretical dreamland that has no relation to reality and that will lead to your being ignored.

As always, there are subtleties to knowing about the observations. Sometimes it is a good idea, before looking at what was reported, to sit down and decide what you expect. If you're right, you have more confidence in your reasons than if you wrote down a post facto rationalization. If you're wrong, you have something to learn, which you wouldn't necessarily have known if you had just read the answer.

In the same spirit as point 5 above, you need to have a sense of what parts of the observations are important or should be modeled. Astronomical sources are finicky, individualistic things, and if you try to explain everything about every source you'll end up barking at the moon while naked. Picking out the relevant parts is an art, which in many cases can be aided by asking simple questions such as "what observational properties are in common among several sources of the class I am studying?". The detail to go into also depends strongly on the state of the theory and the level of the observations. If you're doing asteroseismology theory, you have to work at a pretty fine level of detail because existing theory does a good job and the observations are incredibly detailed and accurate. If you theorize about cosmic rays at energies past the Greisen cutoff, you have a lot more freedom since only a small number of such cosmic rays have been observed and statistical fluctuations probably play a big role.

9. Know your strengths and pick problems corresponding to them

There are many talents that go into successful astrophysics research. These include excellent analytical skills, the ability to write reliable computer code, and a general intuitive physical understanding. Not everyone has all of these, and very few have them all in equal measure. Be honest with yourself about where your strengths and preferences lie, and choose problems correspondingly. I recommend trying out different approaches, because this allows you to strengthen weaknesses and also gives you hands-on experience in multiple different paths.

10. Don't focus on too narrow an area

There is frequently an attitude that graduate school is a time to become expert in one area, and that understanding of a wide variety of subfields in astrophysics is a waste of time. I disagree strongly with this attitude. The possession of breadth is very attractive to potential postdoctoral employers, because it means that you will be able to switch projects with aplomb. You will also be able to talk to a wide range of people and exchange ideas. Study of different fields also allows you to test your intuition in many circumstances, and gives you a chance to borrow techniques that other people may not have thought to apply in your area.

There are a number of ways to enhance your breadth of knowledge in astrophysics once you're out of the classroom. You can read selected Annual Reviews of Astronomy and Astrophysics articles. You can read the abstracts of ApJ Letters to get an idea of what is hot in other fields. You can go to colloquia on subjects disjoint from your own; it helps to do a bit of reading before, in case the speaker doesn't give enough background. You can go around to grad students, postdocs, and professors to ask what they are doing, since you can ask questions if you don't understand. Your ultimate goal is to have deep knowledge in the subject of your thesis, but broad knowledge in other areas.