

How to do research

March 6, 2013 Bill Freeman, CSAIL, MIT
Computer Science and Artificial Intelligence Laboratory
Massachusetts Institute of Technology

The jump from problem sets to research can be hard. We sometimes see students who ace their classes struggle with their research. In little bites, here is what I think is important for succeeding in research as a graduate student.

- The first advice can go on a bumper sticker: “Slow down to speed up”. In classes, the world is rigged. There’s a simple correct answer and the problem is structured to let you come to that answer. You get feedback with the correct answer within a day after you submit anything.

Research is different. No one tells you the right answer, we don’t know if there is a right answer. We don’t know if something doesn’t work because there’s a silly mistake in the program or because a broad set of assumptions is flawed.

How do you deal with that? Take things slowly. Verify your assumptions. Understand the thing, whatever it is—the program, the algorithm, or the proof. As you do experiments, only change one thing at a time, so you know what the outcome of the experiment means.

It may feel like you’re going slowly, but you’ll be making much more progress than if you flail around, trying different things, but not understanding what’s going on.



Figure 1: Research advice for a bumper sticker.

- Please don’t tell me “it doesn’t work”. Of course it doesn’t work. If there’s a single mistake in the chain, the whole thing won’t work, and how could you possibly go through all those steps without making a mistake somewhere? What I want to hear instead is something like, “I’ve narrowed down the problem to step B. Until step A, you can see that it works, because you put in X and you get Y out, as we expect. You can see how it fails here at B. I’ve ruled out W and Z as the cause.”
- “This sounds like hard work.” Yes. It’s no longer about being smart. By now, everyone around you is smart. In graduate school, it’s the *hard workers* who pull ahead. This happens in sports, too. You always read stories about how hard the great players work, being the first ones out to practice, the last ones to leave, etc.
- “How do I get myself to work hard enough to do research well?” It all plays out if you love what you’re doing. You become good at it because you spend time at it and you do that because you enjoy it. So pick something to work on that you can love. If you’re not the type who falls in love with a problem, then just know that working hard is what you have to do to succeed at research.
- I have to note that the above isn’t completely true. Beyond working hard, there’s also *steering*. We’re like boats. We need motors—that’s the working hard part. But we also need a rudder for steering—that’s stepping back periodically to make sure we’re working on the right thing. On the topic of steering, I find time management books to be very helpful. They teach you how to spend your time solving the right problems.
- There’s a concept I want a simple phrase for, and maybe you can help me think up a good name. It’s the simplest toy model that captures the main idea. TSTMTCTMI ? Anyway, simple toy models always help me. With a good one, you can build up intuition about what matters, which is a big advantage in research.

Here's an example. The color constancy problem is to estimate surface reflectance colors when we only get to observe the wavelength-by-wavelength product of the each surface reflectance spectrum and the unknown illuminant spectrum. A toy model for that problem is to try to estimate the scalars a and b from only observing their product, $y = ab$. There's a surprising richness even to this simple problem, and thinking about it allows you to think through loss functions and other aspects of Bayesian decision theory. I co-authored a paper that discusses $y = ab$ for much of the manuscript. Another toy model: as a proxy for complicated shaded surfaces, a single bump. You get the idea. Having the intuitions from working




Figure 2: A toy model for the color constancy problem: $y = ab$

with toy problems gives you a big advantage in the research, because you can figure out what will work by thinking it through with your toy model.

- A parable, as told by my friend Yair Weiss: There is a weak and a strong graduate student. They are both asked by their advisor to try a particular approach to solving a research problem. The weak student does exactly what the advisor has asked. But the advisor's solution fails, and the student reports that failure. The strong student starts doing what the advisor has asked, sees that it doesn't work, looks around within some epsilon ball of the original proposal to find what does work, and reports that solution.




Figure 3: The parable of the two students

- Sometimes it's useful to think that everyone else is an idiot. This lets you do things that no one else is doing. It's best not to be too vocal about that. You can say something like "Oh, I just thought I'd try out this direction".
- It's also sometimes useful to remember that many smart people have worked on this and related problems and written their thoughts and results down in papers. Don't be caught flat-footed with a large body of closely related literature that you aren't familiar with.
- Here's how a business school might talk about your research. You have a brand: you. There are many impressions you want to build up about your brand: that person always does great work, they have good ideas, they give great talks, they write wonderful software. Promote your brand. Build up a great reputation for yourself.
- Cultivate your strengths and play to those strengths. Some possible strengths: being broad; creative; a great implementer; great at doing theory.



Figure 4: Nurture your research brand.

- Please don't report to me, "This instance doesn't work". Why doesn't it work? Why should it work? Is there a simpler case we can make it work? Do you think it's a general issue that affects all problems of this category? Can you think of what's not working? Can you contort things to make an example that does work? At least, can you make it fail worse, so we understand some aspects of the system?
- I love to hear about progress when I meet with students, but note that I have a very general notion of progress. Progress can include: "I've shown why this doesn't work", "I've simplified the task to get it to start working.", or "I spent the whole time reading because I know I have to understand this before I can make any progress."
- Please don't hide from me. Let's talk. I like it when you track me down and insist that we talk, for example, if I've been traveling.
- For a presentation to the visiting admitted MIT EECS graduate students, I emailed all CSAIL researchers and faculty members, "Please send me what you think is the most important quality for success in graduate school" I compiled their responses (along with photos of the responders) into slides that are available online:
<http://people.csail.mit.edu/billf/talks/10minFreeman2013.pdf>
I think it's a lot of good advice about research.
- One final note about doing research: I hope you love it. I certainly do. The research community is a community of people who are passionate about what they do, and we welcome you to it! ¹

¹Photo credits:

1. <http://bethpartin.com/denver-photos-vw-bus-on-capitol-hill>
2. D. H. Brainard and W. T. Freeman, Bayesian Color Constancy, *Journal of the Optical Society of America, A*, 14(7), pp. 1393-1411, July, 1997
3. http://www.123rf.com/photo_2851187_the-two-students-with-the-book.html
4. <http://hms-somerset-co.blogspot.com/2012/12/apple-logo.html>