CSE 501: Introduction to Graduate Study in Computer Science & Engineering Fall 2024

HOW TO DO RESEARCH
Notes for discussion on Tuesday, Sept 10, 2024
(Original notes by Atri Rudra, Updated by Oliver Kennedy)

WARNING: Opinion presented below is mine and is biased. Please note that not all (or even any?) of these might be applicable to (some of) you.

Listed below are some sketchy notes on what I plan to discuss in class on October 7, 2014. The topics are listed in the order of number of votes they received on piazza.

1. (**Industry vs Industry research lab vs academia**) Below are some dimensions under which you can compare the three. (Note: They are listed in my biased order: I start off with things that put academia look in better light. Also when I said academia below I'm referring to a primarily a research job: the demands on a primarily teaching job are pretty different.)

   (a) (FREEDOM OF WHAT YOU WORK ON) There is most freedom on things you can work on in academia (can literally work on anything you want), followed by industrial research labs (can work on anything within the umbrella of things that the lab works on) and finally a development job.[1]

   (b) (FLEXIBILITY OF SCHEDULE) You can typically schedule when you meet people and when you teach in academia. In industry (and maybe to a lesser extent in research labs), scheduling can be much less flexible. (This flexibility can be handful when e.g. your child gets sick and you cannot get a babysitter at a short notice.) And the biggie: in academia you typically have the summer off from non-research duties, while you will not get such a break in industry or industrial research labs.

   (c) (STABILITY) At least so far an academic job is more stable than one in industry (and even in industrial research). On the other hand, the tenure process to get a permanent position can be stressful.

   (d) (STUDENTS) Mentoring a student is a great experience but it is very nerve-wracking. You have the responsibility to make sure your student does well, gets a good job etc. This is great when the student does well. When the student does not do well, then things can get stressful. One has to do mentoring in industry but probably not to the extent as one does for students in academia.

   (e) (PAY) At the beginning there won't be much difference in pay between an industry position and an academic one but over time the difference accrues (it can become a significant multiplicative factor).

   (f) (KIND OF WORK) There are certain work that can only be meaningfully done in the industry (and vice versa). E.g. if you want to really work with big data, then you will find "real" data only in the industry.

[1] Of course if you have your own startup then presumably you are working on something that you want to do. I will be ignoring startups in this discussion.

(g) (NUMBER OF POSITIONS) Industry has lot more openings than academia (as well as industrial research).

(h) (GRANTS) Don't have to worry about this directly in industry but can be a huge time sink in academia (and can impact for practical purposes how much freedom you have in picking topics to work on). In industry you may need to 'sell' your ideas, but it can be easier to make a compelling argument about the bottom line.

(i) (TEACHING DUTIES) Teaching is fun but it still takes time. Does not happen in industry (unless you volunteer to teach somewhere).

(j) (SERVICE) There is a lot of service that one has to do in academia: serve in various committees etc. Can be a huge time sink. Generally nothing comparable happens in industry.

2. (**How to look for a research problem**) There is no set path here. Some things that can work:

   (a) Ask your advisor for a problem to work on. (Ideally, you should only be doing this at the start of your Ph.D– hopefully by the end of it you'll be able to generate problems on your own.)

   (b) Read up papers from recent conferences: good papers will list our open problems. (Warning: the authors are probably already working on those and they have a head start on you.)

   (c) Talk to other students: They might have an interesting open problem for you.

   (d) Attend talks not in your area: you might have the right tools to solve someone else's problem. (This might be more applicable once you have spent some time building your tools during your Ph.D.)

   (e) Keep an Idea Journal: Ideas rarely come to you when you have time to work on them. Keep a logbook (note on your phone, cloud service, etc...) where you jot down any and all ideas that come to you. When you need an idea, revisit the journal.

   (f) Seminars/Reading Groups: We have seminars for at least DB/PL, Theory, and Networks.

3. (**Balance between research and personal life**) Some pointers in no particular order,

   (a) You need a release valve: if possible make it something that is physical. Maybe cook or go exercise. Try and do something that is as far removed from CS as possible.

   (b) If you advisor is a slave driver, make sure you try and set boundaries. Things do get crazy around deadlines and all of you will be working crazy hours but on an average day you should be able to have time to unwind.

   (c) If you have family staying with you, make sure to carve out time for them. (And don't just talk about your problems during the family time — listen to what they have to say.)

   (d) UB-organized off-hours events. We have board game nights, movie nights, hack nights, and more.

4. (**What to do when you get stuck**) Most important thing is to take a break. Do something that is a release valve for you.

   - Maybe think about a different problem for a while and get back to your original problem later.
   - Take a walk / bike ride / etc...
   - 'Rubber duck': Try to explain your problem to someone: A labmate, friend, housemate, etc... or if no one is available, find a rubber duck and explain it to them.

- Write an email to your advisor.

5. (**What happens in the summer?**) For students: work as a (paid or unpaid) RA or get an internship. There are no TA opportunities during summer but there are limited number of summer courses that are taught by Ph.D. students. For faculty: generally free time to do research.

6. (**Smarts vs persistence**) In the long run persistence trumps smarts. In research no one knows the answer so in some sense everyone is stupid. The big results generally come to those who work on a problem for a very long time (multiple years).

7. (**Experience vs smarts**) Do not confuse experience with smarts — when you start off almost everyone who appears really smart does so because they have had a lot of experience thinking about certain kinds of problems.

8. (**Abstraction**) One specific way that experience manifests is in abstraction. When you first approach an idea, a paper, or a talk, you won't understand every concept in the paper exactly. This is *ok!* To quote a wise person on the internet [2]: *I think the single most useful tool for novice learners of subjects, especially subjects with a very wide, very deep, well considered body of knowledge; is to get comfortable with imperfect abstractions. Just get comfortable saying to yourself "so there is this thing called x; and I don't really get it, but it seems to do y. Dunno how it does.".* . The real difference between novices and experts is that experts have a bigger well of experience to draw on with a wider variety of abstractions available, and more nuanced understanding of how and why they differ.

9. (**How many problems should one think about at any given time**) At the very beginning probably only one. But generally I recommend at least two (the upper bound really depends on what you can handle). However, choose diverse problems: both in terms of hardness and probably topics. This helps you when you switch from one problem to another once you get stuck on the first one.

10. (**Communicating with your advisor**) Very crucial to keep talking with your advisor.

    - Bottling up problems does not do anyone any good. If something is bothering you, talk to your advisor about it. Generally, you probably should not tell your advisor all the details in your personal life but if something major is happening that is affecting your work, talk to your advisor about it — they will (or definitely should) understand.
    - Everyone has their own ticks: you should try and understand what makes your advisor tick. At the same time, your advisor will be trying to figure out what makes you tick — so talking about things can only help.

11. (**Quantity vs quality**) Quality always trumps quantity. Having said that by definition it is not possible for everyone to do great work all the time. So it is important to keep working on things and to communicate your result with your research community. Also sometimes the main thing that you get out of working on a paper will not be the paper itself: in two of my major results the main ideas came from other papers that I had worked on earlier (one of earlier works never got published in any conference or journal and the other earlier work probably no one outside of the reviewers read it). Also it is better to get your first paper our sooner rather later.

[2] `https://hackers.town/@rgegriff/108668904927661605`

12. (**Choosing an advisor**) Few points in no particular order:

    (a) (JUNIOR VS SENIOR FACULTY) Typically, junior faculty members will be more engaged (since they probably have fewer students/responsibilities). However, their main aim is to get tenure, so if you are not interested in their pet projects then you probably will not get as much attention as you would otherwise. On the other hand, a senior faculty member is more likely to advise you on a topic that he or she is not that familiar with. On the other hand, it is possible that having tenure might result in some faculty member being less engaged.

    (b) (HANDS ON VS. HANDS OFF) Again depends on what works best for you: some students work better when left alone and some student prefer to get more guidance. Ideally, an advisor should be hands on at the beginning and then be more hands off towards the latter part of your Ph.D. Again, if you have a preference, communicate that with your advisor.

    (c) (ONE ADVISOR VS ANOTHER) If you have more than one potential advisors in mind, go talk to both/all of them. Ideally you should have a small list: then you can start sitting in research group meetings with your potential advisor(s) and/or talking to them. I did something similar at the beginning of my Ph.D. career where I had a list of three faculty members and I worked with each one of them for one semester. (At the end I moved to another school to work with a fourth person.)

    (d) (THE CHOICE ISN'T FINAL) The advisee/advisor relationship is like any other relationship. There's always going to be some level of friction: Your advisor is here to push you to excel. On the other hand, maybe you're not passionate about the topic. Maybe your advisor has different goals than you do. If it's not going to work out, know that breakups do happen. At UB, we'll support you briefly while you look for a new advisor.

13. (**Managing your advisor**[3]) Few points in no particular order:

    - (ASK YOUR ADVISOR) Your advisor is there to help. If you want introductions, or to attend a conference, your advisor won't be able to help you if you don't ask. If you do ask, the worst that will happen is that they say no... but they can't say yes if you don't ask. Ideally, give your advisor plenty of lead time before any deadlines before asking.
    - (FOLLOW UP) Most faculty have very bad short-term memories. If you need to make sure that something happens, make sure to follow up. If you need to meet, follow up with your advisor to make sure the meeting is on their calendar. If you need them to fill out paperwork, follow-up to make sure they're getting to it. If you don't hear a response to a request within 2-3 days, follow-up! (it's not nagging... often I'll thank you for reminding me)
    - (CHECK IN WEEKLY) Even if you have no updates, make sure to meet at least once per week. You may not realize that you're stuck in an avoidable rathole until you actually try to describe your situation. This is especially true if you've spent the last 2-3 weeks on debugging or implementation efforts.
    - (ATTEND GROUP MEETINGS) Every faculty has some form of group meetings or something analogous. This is *really* useful for getting to know your group, and a great way to cross-polenate ideas. It's also a good way to mark time.

[3] `https://greatresearch.org/2013/08/14/managing-your-advisor/`

- (ASK FOR ADVICE) If you feel that you're not getting enough time with your advisor, just ask for more. Different students have different needs, and different phases of research (e.g., just prior to publication) can demand different levels of investment. If you have concerns that aren't being addressed, or questions, don't hesitate... just ask your advisor. Even if they can't talk right away, they'll usually find some time to chat.
- (KEEP ASYNCHRONOUS QUESTIONS SHORT AND FOCUSED) Try to keep questions in asynchronous communications media (email, chat) short and focused: What did you try, what did you expect to happen, what did happen, and what do you think the difficulty is. Keep requests to below 1 paragraph.
- (KEEP MEETINGS FOCUSED) Spend a bit of time before meetings sanity checking your results. If you're working with a graph (or numbers that could be graphed), ask yourself what you would expect the graph to show. Does the graph actually show it? What is the high-level point you're trying to show? Prepare by asking yourself what you would like advice on during the meeting, and try to focus the discussion on that.
- (TAKE NOTES) Keep a lab notebook (or e-book). Don't record *everything* that happens during the meeting, but note down highlights. Specifically jot down major points that were made, and anything that you (or your advisor) need to do before the next meeting ('action items'). Follow up with your advisor to make sure they're doing the thing.
- (FORMALIZE EARLY AND OFTEN) Paper-writing is deeply part of the research process. The data, definitions, and/or theorems tell a self-coherent story, and understanding where that story is going, is critical for doing good research. Draft a (bad and replaceable) introduction to your paper as soon as you have some idea of where you're going. As you come up with ideas, put them into LaTeX. Don't just write down the ideas, but try to introduce them: What problem does the idea solve? Why is it interesting.