Advice to senior grad students and postdocs

These notes are an adapted summary of the contents of 4 professional development panels for young scholars, held as part of the American Mathematical Society (AMS) Mathematics Research Community (MRC) on Network Science, in June 2014. The focus of the MRC was on mathematics and computer science, and especially network science; however, many of the comments generalize to other fields.

Panelists were the MRC organizers, Profs. Mason Porter (Oxford), David Kempe (Southern California), and Aaron Clauset (Colorado), along with an ad hoc assortment of individuals from the MRC itself, as per their expertise. The notes were compiled by MRC participants, and then edited and expanded upon by Aaron Clauset (© 2014) for clarity and completeness.

1. The Academic Job Market

Opening remarks: The faculty hiring process is much more personal than the process to get into grad school. Those who are interviewing you are evaluating whether you should be a coworker for the rest of their careers! The single most-important thing about preparing to apply for faculty jobs is to have a strong CV for the type of job you're applying for. If you're on the tenure-track, that nearly always means being able to show good research productivity for your field (publications) *and* having them be published in the right places for your field.

Q: *Where did you find job postings? Where did you search?*

A: It depends on the type of job and the field. For math: AMS weekly mailings, back of SIAM news. For physics: the back of Physics Today. For computer science: cra.org/jobs, Communications of the ACM. For liberal arts colleges: chronicle vitae. In general: mathjobs.org, academicjobs.org, and ask your supervisor(s) or coauthors.

Q: *When do you apply?*

A: The U.S. market is, for the most part, seasonal. The seasonality differs by field. Biology searches may start in September, with interviews in November and December. Math and computer science tend to have applications due in November, December, and maybe even January. In the U.K., institutions tend to hire whenever, regardless of season. Timing for interdisciplinary positions may be a little strange. It is worth figuring out 6 months ahead of time what the usual timeline is for your field.

Q: *What kind of department should you apply to?*

A: If you're in department X, you will be expected to teach courses in department X. (At most institutions, you will teach a mixture of undergraduate and graduate-level courses, but not always within your research speciality.) It may be better to have your publications match the departments to which you apply; for instance, if you're interested in jobs in math departments, you should be publishing in the SIAM journals. You should also get letter writers in that field, since their name will be more recognizable to the hiring committee (and thus carry more weight).

Q: *What should you put in a cover letter?*

A: The cover letter is the first (and sometimes only) thing the hiring committee sees. In some fields, the cover letter is 1 full page of text and serves as a complete abstract of your application packet (i.e., it describes your preparation, major research areas and achievements, and intended future research agenda). If you have a specific interest in a department / location, say it in the cover letter (e.g., "I have family living in X and want to be close to them") since this signals to the hiring committee that you're genuinely interested in their institution. Also, mention the people in the department whom you would like to look at your application. Mention a few specific things about the individual advertisement (no one likes to feel spammed). Finally, the cover letter is your one chance to explain anything that might look like a red flag to the committee.

Q: *What should you put in a teaching statement?*

A: At research universities, teaching statements are usually the last thing that is read. For junior-level positions, their contents often cannot help your changes, but a bad statement can hurt them. At liberal arts / teaching colleges, a compelling teaching statement is very important.

Q: *What about letters of recommendation?*

A: Letters of recommendation are the second most important thing in your packet (the most important being your publication record). The best letters are those that can state firmly that you are in the top whatever percent of students or postdocs. Their description of you is the most important, and their own fame is second. There are some cultural differences between the U.K., U.S., and other places in terms of how glowing they will be. An excited letter from an unknown writer is more important than a mediocre letter from a famous person. The absence of a letter from a PhD or postdoc advisor will be interpreted as a red flag.

Q: *Are software and blogs good or bad?*

A: Sometimes good, sometimes not. Don't do these things at the cost of your own research. If you have specific reasons for doing these things, emphasize them in your research statement as "sweeteners" to your strong publication record. For tenure-track faculty jobs, these things generally cannot compensate for a poor or mediocre publication record. The research itself is the most important thing.

Q: *How does the hiring committee work?*

A: At most institutions today, the ratio of candidates to faculty jobs is roughly 100:1. At major research institutions, about 60% of those candidates are not competitive to begin with; it's the other 40% you have to beat. This means the hiring committee has an enormous job to do just to narrow the pool to the 10–20% or so that they'll scrutinize closely. Your goal is to make it into that group, so that your star qualities can be properly noticed.

A common strategy that hiring committees take is to progressively pay more attention to a progressively smaller pool. Your goal is to get through the first few waves of filtering until you can get a serious look by the committee. Two very common reasons a candidate is dropped from the pool during these early evaluations are (i) their area of research is not a good match to the areas of interest in the search (including not looking like the kind of candidate the committee thinks they want, e.g., because their work appears in unusual places), and (ii) their research productivity is not good enough (usually controlled for time since PhD). Both are subjective criteria and vary by search. In general, the more prestigious your PhD, the more prestigious your publication venues, and the more prestigious your letter writers, the better you will fare.

Q: *What about the interview itself?*

A: Usually 1-2 days of intense, back-to-back meetings with faculty, plus a meeting with a group of graduate students, plus 1-2 dinners with faculty, plus a job talk (about your research), and sometimes also a "teaching talk." In your job talk, you need to convince them of why they should hire someone doing exactly the research you're doing. Make the audience excited. Make it related to things they know about. Be sure to look at the webpage of every person that might be in the room. Be sure to ask for your meeting schedule in advance, and then read up a little about each person you will meet.

Pro Tips:

- "Exploding offers" (offers that expire after a few weeks) may be used by lower-tier institutions when trying to hire a person likely to have offers from higher-tier institutions. But, deadlines are often negotiable. Play nice. It's often not malicious, but rather just to proceed quickly down the ranked list of candidates. Moreover, if you turn it down in a friendly conversation, you may be able to negotiate a "if you are still interested in me in a month, please let me know."
- During the year before you apply, figure out what departments you'll be applying to, and be sure to have some publications and talks at major conferences for that type of field or department.
- Don't pad your CV. Put all preprint and in-prep publications in a separate, clearly-labelled section. CV readers will look at your PhD, your research interests, and then your publications. Awards (e.g. Best Paper) and high quality venues are more important than quantity.
- You could email people at the target department(s) saying "Here's a paper, btw: I'll be applying soon." If you're uncomfortable with that, your advisor could do it.
- If you are applying to a lower tier school than your pedigree, tailor the application well. You must send a very strong signal that you are serious. (Otherwise, they may not even interview you.)

2. Life / Work Balance

Opening remarks: "Academia is like art because we're all a little crazy."

Productivity often scales with time spent. A good strategy is to find enough of a balance so that you don't implode or burn out or become bitter. The best way to find that balance is to experiment! Social norms in academia are slowly shifting to be more sensitive about work/life balance issues, but academia changes slowly and sometimes you will feel judged. Often, those judging are senior faculty, possibly because of classical gender roles in the family and the fact that their children (if any) are usually grown. Telling people you're unavailable is uncomfortable, but you will get used to it. Pressure will be constant, so if you want a life and/or a family, you just have to do it. Routines can be powerful—make some rules about when your non-work hours are during the week and stick to them.

Q: *Having children?*

A: Most institutions have a standard paternity/maternity leave option: one semester off of research/teaching/service plus a one-year pause on your tenure clock. If you think you will have children while being faculty, ask about the parental leave policy during your job interview.

Faculty with small children often have to deal with scheduling constraints driven by day care hours, or at-home responsibilities for child care; they are often simply unavailable nights and evenings, so be sensitive to that (don't assume they will be available for work stuff then).

Juggling a brand new faculty job *and* a new baby in the same year can be done, but it can also burn you out.

Q: *Burnout?*

A: It's hard to get numbers on burnout rate, in part because there are varying degrees of "burnout" and different people burn out in different ways. Most tenured faculty are not completely burned out; true burnout often turns into leaving academia. On the other hand, some faculty have real breakdowns and then get back on the horse. Other faculty give up on the "rat race" of fundraising and publishing in highly competitive venues and instead focus on teaching or service. There are many ways to stop being productive and lose the passion.

One strategy is to promise yourself that once it stops being fun, leave and go get a satisfying 9-5 job (that pays better).

Q: *Service?*

A: Service (to your department, to your university, and to your research community) is an important part of being a professor. You will get asked to do *many* things, many of which you've never done before, some of which will sound exciting. As an early-career person, you should learn to say "no" to things and feel comfortable with that decision. Until you have tenure, it's okay to be fairly selfish about your service—think about whether saying "yes" will have a benefit to your own research efforts. If the benefit is marginal, then you should probably say no.

There are a lot of factors that go into whether or not you say yes to something. It's important to learn to tell the difference between something you should say no or yes to. A key part of this is having one or more senior faculty mentors you can ask. Ideally, have one inside your department *and* one outside your department but within your research community.

Q: *Summers?*

A: If you're willing to set yourself up for it, then you can readily take a month-long vacation with absolutely no contact. Tell your department head that you're not bringing your laptop. That being said, summer is often the time where many faculty try to focus exclusively on research, since they're not teaching. At most institutions, it's normal for regular departmental committees to not meet, so you often get a break from your departmental service obligations then, too.

Q: *Work hours?*

A: How much you work each week is really up to you. Some people work 80-85 hours during terms, and 70 between terms. A common number kicked around is 60, and relatively few people work a lot less than that. For the most part, faculty work these hours by choice. The great advantage of faculty life is that your schedule is pretty flexible, which allows you to carve out specific time for other things (e.g., life / family). Many faculty work 9-5 on campus, and then add other hours at home or otherwise off campus. Some others work long hours during the week and then are offline on the weekends.

Q: *Showing up to the office?*

A: If you don't get "face time" with your institution and the people evaluating your tenure case, then they will form negative opinions about you. So go into work often. And, spend time "in your lab," with your students. Good idea to have lunch with every one of your fellow tenure-track faculty during your early faculty career.

Q: *The two body problem?*

A: Solving the two-body problem (marriage with another academic or other professional career type) can be tricky. Start talking about it with your partner long before you start applying to jobs. One solution: make a list and let your partner cross off the things that don't make sense. In job negotiations, there are things that the department can do, such as interview/hire your spouse (or encourage/fund another department to do so). If your partner is not an academic there are few things the university can do, but often the more senior people have contacts and that can help.

One strategy is to always go for the interview, get the offer first, and think about it later. Departments often want to know ahead of time whether they'll need to help with the two-body problem in order to get you to say yes. (But, they are legally not allowed to ask you if you have a partner, so you have to bring it up.) This can (but not necessarily) hurt your offer. Also, when women interview, they get assumptions imposed on them, such as the existence of a two-body problem. Some women don't wear a wedding ring to an interview in order to avoid those assumptions. One possibility is to consider saying something in advance along the lines of "my husband is excited and there's no problem.

Q: *Travel?*

A: Many strategies. Mostly depends on your personal preferences. A popular strategy is to travel no more than once a month. Also consider picking trips on which you can bring your family and/or do some extra traveling. As a junior person, however, traveling is in part about reputation-building, and is a necessary part of academic success.

3. Interdisciplinary Research

Opening remarks: Sometimes, the most interesting problems come from interdisciplinary fields, and interdisciplinary researchers are becoming more and more common. As network scientists, we tend to fit in with many disciplines. That said, the most important thing you have is time; therefore, choose your collaborations wisely. Interdisciplinary work can be divided into collaboration and publication, and each of these has its own set of difficulties. A common experience with interdisciplinary work is this:

Any paper that aims for the union of two fields will appeal mainly to the intersection. — Jon Kleinberg

Q: *What's the deal with interdisciplinary collaborations? How do they impact your academic reputation?*

A: There are three main points to consider when choosing interdisciplinary collaborations, and how they impact perceptions of your academic reputation.

First, academia is very tribal, and the opinions of these tribes with regards to your work can have a huge impact on your career. Some departments won't value work outside their scope. (Some even have a short list of sanctioned publication venues, with work outside these venues counting literally as zero for your assessments.) Other departments are more open minded. In general, it's important to signal to your hopefully-future-colleagues that you are "one of them." This can mean publishing in certain places, or working on certain classes of problems, or using certain language in your work, etc. If you value interdisciplinary work, then you want to end up in a department that also values it.

Second, it's strategically advantageous to be "the person who is the expert on X," where X might be algorithms or statistics or models for networks, or whatever. Your research specialty won't necessarily align completely with any particular department, but it should align well with a particular external research community. In the long run, it is much more important to fit into your community than to fit into your department, research-wise. This community will be the group of people who review your papers, who write your external letters when you go up for tenure, who review your grant proposals, who hire your students as postdocs, etc. The worst possible situation is to be community-less. You don't have to choose your community now, but it helps to choose far enough ahead of your tenure case that you have time to build a strong reputation with them.

Third, make sure the research is interesting *to you.* If your contribution in some interdisciplinary collaboration is to point out that an off-the-shelf algorithm solves the problem at hand, it's probably not interesting to you, even if it's very interesting to the collaborator. Even if it gives you an easy publication, it won't have much value to your reputation

in your community. Your work will be compared to the work of people who do only one type of research in both fields, and might not look particularly good to *any* field.

Be very careful about potentially complicated collaborations in the early stages of your career. Be noncommittal until you're sure that your personalities and tastes in problems match. (Getting "divorced" from a collaborator, once a project has started, can be exhausting and complicated.) Being able to recognize cultural differences is an important first step to good collaborations, and moving forward effectively. Don't burn bridges, but don't fall into the trap of saying yes to too many things. Be open to writing for an audience that is not your primary research community, and be open to learning what makes an interesting question and a satisfying answer in another field.

Q: *What's the deal with publishing interdisciplinary work? Where should it go?*

A: As a mathematical or computer or data scientist doing work in a domain, be sure to engage with that domain's community. This helps ensure that you're doing relevant good work, and not reinventing wheels. Attend talks at other departments at your university, attend workshops/conferences in the domain, and discuss your results with people in the domain audience.

When writing, vocabulary is important. Knowing how to speak another discipline's language will help you write in a way that satisfies reviewers from that community. Less cynically, it also helps the audience of that journal understand your results, which is the real goal. If publishing in the arena of a collaborator, trust your collaborator on the language/writing style.

In general, know what part of the paper is the most interesting, e.g., the mathematics, or the method or algorithm, or the application and relationship to scientific hypotheses, etc., and send the paper to a venue that primarily values that thing. This can sometimes be difficult, since academic tribes are, by their nature, fairly conservative, and attempting to publish a new or interdisciplinary idea can meet with knee-jerk resistance. Interdisciplinary journals like PLOS ONE, which try not to consider domain, can be an okay solution for early work that has trouble finding a home. But, don't overuse these venues, since they tend also to not have a community of readers built in the way regular venues do.

Note: When you interview for a faculty position, among the many questions that you should be asking the interviewing department: "In practice, how is your department interdisciplinary? How do you consider interdisciplinary work when evaluating young faculty (e.g., at tenure time)?"

4. Grants and Fundraising

Opening remarks: In general, only around 10% of grant proposals are successful. But, roughly 60% of submitted proposals are crap. Your competition for getting funded is the non-crappy 40%. Therefore, work hard to polish your proposals, and take as much time as you would a serious or flagship paper. Get feedback from colleagues on your proposals before submitting, and try as hard as possible to get that feedback at least one month before the deadline. (Many institutions have these "mock panels" available, and they are incredibly useful, especially for early career scientists.) Practice makes the master, so consider writing a grant proposal as a postdoc. Having some success as a postdoc will also make you look more attractive as a faculty candidate. Know when the annual deadlines are for the regular grant competitions, and plan ahead. Try to avoid the last-minute crush of writing proposals in two weeks or less.

Q: *What should be in a proposal?*

A: Really exciting research. But, try to propose to do more than just really exciting research. Consider organizing workshops, creating new classes, creating notes, giving public lectures, hosting undergraduates, working with underrepresented groups, running a podcast series, and even teaching in a local high school.

Q: *What kinds of proposals should an early-career person write?*

A: In your first few years as faculty, apply to all the early-career fellowships and competitions that you can comfortably manage. That includes the Sloan, McDonnell, Packard, etc., along with the NSF CAREER award, and the various "early investigator" competitions at the DoD and other places. Figure out what people do in your field and do that too. These awards are sometimes for sizable amounts of funding, but even if they are not, they are often very prestigious.

Q: *How many grants do I need?*

A: This depends on the size of your preferred research group. Many faculty try to keep 2–3 active grants at once, and write approximately 1–2 new proposals per year. As a rough calculation, a "normal sized" grant from many parts of NSF will support 1 graduate student for its duration (plus modest summer salary, travel, and computing equipment).

Q: *Can I propose work that I have already partially completed?*

A: Yes. This is common, and often even recommended. "Preliminary results" make a proposal sound less risky, and basically the reviewers are looking for proposals that are exciting, will advance the state-of-the-art, well written, and exceedingly likely to succeed. If you've already worked out many of the details of the work itself, it is much easier to write a compelling proposal.

Q: *Proposals are often required to be understandable by a broad audience but also include technical details, so how do you balance these requirements?*

A: An advanced undergraduate should be able to understand your proposal with some training. Most panels have some experts who can judge the technical details. A good strategy for learning how to balance technical material versus accessibility is to read other people's proposals, especially successful ones, even outside your field. The first pages of any proposal should be more broadly understandable, while the final pages may be decodable by experts only.

Q: *Can you reuse the same material for multiple grants?*

A: It's best not to double dip. If a grant is rejected, you can usually resubmit it, often to the same agency (although sometimes not more than once). Because you have some feedback and you have already written the first proposal, it's often less work to revise and resubmit a rejected proposal. (But, the goal posts may move with the resubmission because the review panel may be composed of different people with different opinions, e.g., at NSF.) Small amounts of overlap are usually okay, but if you don't have anything new to propose, don't submit a proposal.

Pro tips:

- Calls For Proposals (CFPs) are often difficult to decode, so don't hesitate to ask for help to translate, either from your colleagues or from the cognizant program officer. Usually, the specific words and pitch of a program have been shaped by other researchers' interests, and knowing what those words really mean can help in deciding if your ideas are a good match for the program.
- Proposals are reviewed differently depending on the agency. NSF proposals are reviewed by ad hoc committees of practicing scientists (drawn essentially at random from a particular broad domain). NIH proposals are reviewed by study panels whose membership is fairly stable over time. DoD proposals are reviewed internally, but sometimes with input from outside individuals (who may or may not be academics).
- Don't write the budget yourself. Use the resources of your department. You will eventually learn many things about budgeting, but your time is better spent writing about the science. That being said, you will need to think about budgets a lot because they are what pay for the research to get done (and universities and funding agencies really love to treat them like immutable, sacred documents). Familiarize yourself with the actual expenses associated with your kind of research, and with the projects that you currently and aim to do in the future.
- For NSF, don't budget for funding to support undergraduates during the summer; instead, assume that you will apply for (and receive) an REU Supplement to your award to cover them. The funding rate for these is well above 50%.
- NSF (and some other agencies) have byzantine rules about the structure, format, and set of documents included in a proposal. They now routinely reject without review proposals that don't follow these rules to the letter. Don't be one of those people.
- Ending up with leftover money is not good. Write an accurate budget and spend it. Many agencies (e.g., NSF and NIH) will allow you to do a 1-year "no cost extension" to spend the remaining money.
- Program officers at NSF are typically professors, on leave for 2-3 years, so speak to them at conferences. Program officers at DoD agencies and private foundations are typically professionals (not academics). NSF program officers exert fairly little influence over the review and scoring process of proposals. DoD and foundation program officers exert enormous influence over their process.