

Thoughts on applying for faculty positions

By Arjun Raj, upon discussions with postdocs from around the world navigating the academic job market

These thoughts summarize some of the points that repeatedly came up during my—highly enjoyable!—discussions with postdocs navigating the academic job market for faculty positions. They are also very much a reflection of my own opinions, values, and biases, many of which I know are not shared by my colleagues 😊, so take them as just another datapoint. Also keep in mind that the target job is typically a primarily research oriented tenure track faculty position at a basic science or engineering department in a medical school or engineering/biology department.

Mostly, this doc is meant as a reference with most of my generic thoughts on the process. **However, I did consistently see one critical issue for almost everyone I talked to: when asked to describe their work, the motivation for the work did not come through.** Like, why did you do what you did, and why should I care? I have outlined some thoughts on how to work on this problem in the [research seminar](#) section. If you're going to read any one part of this doc, start there.

[The written application](#)

[Publications](#)

[CV](#)

[Internet presence](#)

[Research statement](#)

[Teaching statement](#)

[DEI statement](#)

[Timing](#)

[Where to apply](#)

[Zoom interview](#)

[In-person interview](#)

[Research seminar](#)

[Chalk talk](#)

[Meetings with faculty and students](#)

[Aftermath](#)

[Negotiating an offer](#)

The written application

A lot of people ask about what one can do to best differentiate their application from all the rest. My advice is plain and simple: **do your best science and publish it as best you can.** Repeat after me: **do your best science and publish it as best you can.** It is the dominant factor in your application. You're probably thinking "well duh, I'm doing that already, so how can I act on this advice?", but as with many things here, I think this

advice matters most in terms of what you spend your time on. Should you spend more time polishing your research statement? Maybe... but better to use that time to *do more science and publish it*.

Publications

The most critical thing after doing the science is to publish it. Why, you might ask? From the perspective of hiring committees, it allows them to read the work to judge its quality. Sometimes you might see an awesome talk, but without the written details, it can be hard to gauge the rigor of the work. Publications also offer a window into how a candidate thinks, especially publications over time across multiple labs. The formal publication is also important because it is an important signifier for those early years of your professorship, helping you to attract students and win early career awards.

Here are some concrete tips:

- **Post your work as a Biorxiv preprint before or when you submit for formal publication.** Doing so enables the hiring committee to easily read your latest work even if not yet formally published. Some hiring committees may not weigh it as heavily for various reasons, but it doesn't ever hurt, and I can say that I have definitely read Biorxiv preprints from faculty candidates.
- **Publish as "high impact" as you can, but don't be a slave to impact.** Sometimes you can waste valuable years chasing after the fancy Nature/Cell/Science paper (and ultimately being unsuccessful), but instead could have used that time to maybe get an additional paper out in a high-quality venue. There is of course a narrative that a "top 3" paper is required to get a faculty position, but hopefully that myth has been dispelled. If you don't believe me, take a look at the records of who gets hired at various places. Yes, many do have papers in the big 3, but many do not. What you also do not see is that many applicants have papers in the big 3 journals but don't make the cut. Again, just do your best science. That's what matters most.

CV

This document is the most widely circulated, so is very important as a concise summary of who you are as a scientist. The main elements are probably known to you: education, papers, awards/grants, leadership and outreach. Here's a few thoughts on this document:

- **I think giving a 1-2 sentence statement of who you are at the top is helpful.** Like "My research program is aimed at discovering the principles of adaptive behavior during oncogenesis by applying both experimental and computational tools to build quantitative models of cellular function." Or whatever. Just a quick, high-level orientation. A busy dean might get this CV forwarded to them, imagine yourself in their shoes.
- **List papers chronologically, newest at top.** Review committees want to see your most recent first author papers foremost.
- **Clearly indicate which papers are reviews and opinion pieces.** Probably best to keep them in a separate section.
- **If you list manuscripts in preparation, keep those in a separate section as well.**
 - Never say "in prep for submission to XYZ fancy journal". That is lame, and some people will be really turned off by that.
- **Include preprints.** Can be either in the list of papers (again, at the top), or a separate section.
- **Do not change the order of names for any co-first authorships from the actual order on the publication.**
- **For major first author papers from your PhD and postdoc, you could consider including 1-3 sentences about why the paper is important.** I think these short summaries are helpful because CVs are often circulated without other documents like cover letter and research statement, and so these

mini summaries help highlight your work. Other colleagues I've talked to feel like it's redundant with the research statement and just makes extra work for candidates. Up to you.

Internet presence

In this day and age, there is simply no excuse for not having some sort of internet presence. Committee members will often Google people, and so you want to have something there for them to click on.

- **Make a Google Scholar page for yourself.** Takes minutes and is super helpful. Just do it.
- **Make a website for yourself.** A simple website is an opportunity for you to provide succinct summaries of your work and what it's about. You can also provide some personal history about your career to date. Keep it short and sweet. I like it if the publications are directly listed on the website rather than linking to Pubmed or Google Scholar, but keep it up to date please :)

Research statement

The research statement is your opportunity to provide a narrative about what your research vision is, both what you have done and (mostly) where you want to go. Thoughts:

- No more than half a page on prior accomplishments. They have your CV. Make it a narrative, too, highlighting your unique qualities.
- Be sure to set up the problem, discuss why nobody has been able to do it before, and why you are uniquely well positioned to solve it.
- Don't be afraid to describe at least 1-2 crazy ideas.

Teaching statement

The teaching statement documents your experiences as an educator and the principles that guide your teaching and mentorship. Depending on the type of position you are applying to, the teaching demands may be quite different, and you may want to highlight different aspects of your prior experience.

DEI statement

With the increased focus on the challenges we face making academia more diverse, equitable and inclusive, many departments are asking for statements from candidates that speak to these issues. In this document, you can discuss your own personal experiences and journey, your ideals and values, and describe any programs you have participated in or hope to develop upon starting your position. Be assured these are discussed and strong ones are noticed.

Timing

Knowing when to apply is tricky and highly individual. That said, there are a few general tips. The most important thing is to wait until your big paper is basically done. Hiring committees are looking for you to publish your primary work from your postdoc. It's important because having that publication when you start on faculty makes it easier to get trainees and especially new investigator grants. The most common mistake I've seen is to apply too early, well before the paper is written. Often, it's because you feel like you're ready to go before the work is actually published. Resist this urge. If your paper is submitted around the summer/fall, you could be okay, but hiring committees know just as well as you do that it takes forever to publish these days.

If you're going to ignore this advice and go on the market anyway (and I know many will), then I would say the most important thing is to only apply to a small number of places, ones that are really at the top of your list. If it works out, then great. If not, then you can try again next time with a stronger application and a broader search. That way, you won't wonder "what if" if you take a less exciting job early. Also, keep in mind that interviewing and applying takes a lot of time, and it can be really hard to do science during that time. Thus, there is a significant opportunity cost to applying twice.

It is also important not to apply too late. Once your paper is out for a while, people start to wonder if you have another paper coming. Now you've just unwittingly put more pressure on yourself. So once your major paper is ready, don't wait. Also, I personally would not wait for a grant. If you have the paper but no grant, go for it. I don't think having a grant matters nearly as much as people think.

Where to apply

Where to apply is highly individual and depends on your own values and goals, both personal and professional. Places can have very different atmospheres, trainee quality and other things that can greatly affect your professional quality of life. Med school vs. undergrad institution is also a big difference, with hard money and teaching at the latter being the biggest difference. Also, keep in mind you might be at this place for a long time in your life. Up until now, you've probably put up with some amount of personal inconvenience for your professional goals because positions are temporary. I'm not here to tell anyone how to make these sorts of calls for themselves, but just keep in mind that you might be at this place for a long time.

Whatever your dimensions, you probably have a list of places you would be willing to go to and a line below which you wouldn't. Do not explore the latter. Make a list ahead of time of places and their relative ranking for you so that it's there for you. When you're actually in the middle of the interview process, you will go through a lot of ups and downs as some things seem to work out and others don't. It can be really hard to remain objective about some of your inherent values and so forth. So just refer back to that. I'm not saying that you can't revise your opinion about a place over time and through the visit, just saying that you can end up in some really weird headspaces during the process.

Zoom interview

Zoom interviews are a relatively new aspect of the faculty job search, one that has been accelerated dramatically by the pandemic. Search committees will now often have a shortlist of candidates of interest that they will interview by zoom for 30 minutes or so, then move forward with some of these for in person visits. Overall, I think they save valuable time for all involved.

The format of these zoom interviews vary dramatically and I can't say I have any comprehensive knowledge of all such types. I think they typically consist of a very short talk about your research followed by a question and answer period. The question and answer period can either have canned, preplanned questions (you can probably guess what these might be, like "how will you distinguish yourself from your mentor" type of stuff) and maybe some freeform discussion.

My primary advice for these is the same as for the main research talk (more extensively described below), which is to present your science with clarity and contrast. **Emphasize how the field was before you showed**

up, what your work opened up, and why you are the right person for the job. You can spend a little time on future directions, but in my opinion, in the best job talks, the future directions are obvious. Packing all this information into 10-15 minutes is not easy, but that is your job. Do it well and you will be rewarded.

In-person interview

If you got an interview: congratulations! Having witnessed how many applications get reviewed each cycle, it really is something to make it to the interview round. (Also if you don't make the interview round, it's okay. There could be a lot of reasons for it that don't have anything to do with you or your science.)

The typical interview consists of a 1-2 day visit at the host institution. You will meet with a bunch of faculty for 1-1 meetings, give a research seminar, give a chalk talk, sometimes meet with students, and have dinner with a few faculty.

One key general tip: you will have a point of contact in the department. Use that contact. Ask questions if you have any. Sometimes your point of contact will be the person who was the "champion" for your application, and can provide some very useful guidance.

Interviews are really anxiety producing because of the great deal of uncertainty you face: what are these people like? Will I screw up something? What the heck even is a chalk talk? My number one bit of advice for the process as a whole is to just be yourself. A lot of unexpected things might happen, and you will be best represented by handling them with grace, respect, and a positive attitude. As for preparation, there are a few concrete things you can work on, the most important being the...

Research seminar

To me, the single most important part of the interview is the research seminar. It is your opportunity to share your scientific output, aesthetic, and vision with the community. In my opinion, it is by far the most influential aspect of your recruitment and worth spending time to improve. In my experience, while there are a lot of other parts of a visit (1-1 meetings, chalk talk, etc.), the research seminar becomes the lens through which the rest of these interactions are viewed. Give a good research seminar, and it will positively color how your chalk talk goes. Give a bad one, and you have a lot of ground to make up.

The irony is that while the research seminar is (in my opinion) the most important part, I think it's the one that often gets the least attention from candidates. They tend to think "Hey, I've given tons of seminars, I know how to do that, but... I don't know how to do a chalk talk, so I'll spend all my time on that". Big mistake. The vast majority of research talks have a lot of... ahem, room for improvement, and yours is likely no different. How? [Here are some tips on how to improve your talk.](#)

In my discussions with applicants from many disciplines, I've noticed many make the following critical mistake in describing their work: **they fail to set the stage for their story.** Like, in the 5 minute version of their spiel, they say something rather generic, like (completely made up):

"I study single cell epigenetics. I made a tool for measuring methylation in single cells. Using this tool, I analyzed a cancer cell line, showing that individual cells vary in their methylation patterns. These patterns correlate with various changes in therapy resistance."

Or just some statement of facts, like:

“I study neuronal growth. I showed that the factor XYZ76 is responsible for the growth of axons towards their targets with an in vivo CRISPR screen. We validated the hits in our in vitro model.”

Sounds fine. But several questions remain unanswered:

1. What is the problem and why is it worth solving?
2. Why didn't anyone solve the problem before?
3. What was your unique insight that led to your solution?
4. What, specifically, did you learn?
5. Where will this lead? [Often implicit]

Fundamentally, these descriptions provide no context for what the current state of the field is and no contrast between what the state is and what you have done. **Building contrast, in a narrative form, is what makes people get excited about your work.**

Instead, try something like this:

“The process by which cells remember their prior history is important for processes ranging from development to disease. Yet, it has primarily been measured at the level of populations of cells, making it impossible to know how the individual cells responsible for, say, therapy resistance in cancer change their epigenetic characteristics. The technical challenge is that each cell has just two to four copies of the DNA, making it difficult to detect these changes. To solve this, we employed [clever trick goes here]. Using that, we revealed how cells destined to be therapy resistant have methylation at promoters of particular genes. This result points to methylation of these genes as a natural source of variability that could be at work in a number of biological contexts.”

Or

“Neurons must connect to each other in order to communicate. Mostly, the field thinks that a particular factor, SMH99, is responsible for neurons connecting, with the prevailing model being that it acts as a proximity sensor to tell a neuron to stop growing once it hits another one. However, this model doesn't explain how a neuron knows where to go in the first place. We used a new model system with [the specific clever feature] that enabled us to perform genetic screening for directionality separate from proximity. We thus identified a new factor, FML76, that regulates directionality. The mechanism by which it performs this regulation is XYZ.”

These examples are ones I made up really quickly, and so could certainly be improved, but you get the point. If it helps, a guiding principle is the Pixar storytelling formula:

Once upon a time there was _____. Every day, _____. One day _____. Because of that, _____. Because of that, _____. Until finally _____.

Be sure you fill in those blanks every time you talk to someone about your work. And you will talk about your work, over and over and over again. Not every faculty member is going to be able to make it to your research seminar. Not everyone at your chalk talk will have been to your seminar. **You need to have a 1 minute, 2 minute, and 5 minute version of your talk ready to go.**

Keeping the above in mind: PRACTICE. Yes yes, you say, I will. And then you won't. So let me say it again: PRACTICE. Ideally, in front of colleagues who don't know your work very well, so that they can provide an outside perspective. A few notes about getting feedback: people often will not give you good feedback, or may not even have the tools to do so because they can't quite articulate why something is good or bad, or even

recognize that something is bad because all they've seen is bad. For that, some sort of A/B testing can be very helpful. Try a few different things (in particular, implement the above advice and ask for opinions). And practice the short versions as well. No excuse not to, it by definition is short!

Chalk talk

Ah yes, the dreaded chalk talk. Bane of faculty candidates the world over. What should I expect? How should I prepare? What even is a chalk talk? In my opinion, the biggest problem with the chalk talk (for candidates) is that they spend too much time preparing for it when they should have spent that time making their primary research seminar better (see above). I also think chalk talks are kinda useless for evaluative purposes (but I think I'm largely on my own there, so they are probably here to stay).

Anyway...

Chalk talks are basically another chance to interact with the candidate, but behind closed doors and just with faculty (sometimes postdocs are invited). The format is quite heterogeneous, and I can't claim to have seen all possible formats, but over time, it has evolved into something like a "Here's my first R01" with a bunch of Q&A. Format could be actual chalk or slides. The successful ones, in my opinion, just serve to confirm the positive impression left by the primary research seminar. Bad ones tend to expose some degree of shallow thinking, although I find it sort of unfair to judge people on their ability to "think on their feet" given that that skill has very little bearing on the actual job in question, in which you can take your time to think about virtually everything you're actually judged on (papers, grants, seminars). Did I mention I think chalk talks are stupid?

A few tips:

- Sound confident, but accept and engage with criticism. People will ask questions, and some may be tough questions. If they point out a flaw, accept that and see if you can find a solution. Think of it as a conversation between future colleagues. Avoid being defensive about your ideas. Sometimes that can lead to the "blood in the water" effect.
- Chalk vs. Slides: ask your point of contact well in advance. You may have either option. Which you choose depends on preference. The main difference is that chalk talks go a lot slower with actual chalk than with slides. I generally think that's a good thing, since it allows people to digest the material more fully and is a lot more interactive.
- The chalk talk format has evolved from "who knows what we are doing here" to "show me an R01". So it may make sense to talk about specific aims. (On a personal note, I think spending valuable time with colleagues about how your work would fit into an R01 is just about the least imaginative thing ever. Sigh.) I do think it's also useful to outline some more crazy ideas as well to demonstrate your range.

Meetings with faculty and students

The rest of your visit typically consists of meeting with various faculty. Also, sometimes you will meet with students. **All of these meetings are part of the interview.** As such, you are always "on". Again, treat every single person with respect and enthusiasm, and good things will happen. This maxim includes staff. I have definitely heard of candidates treating administrative staff poorly, and that's a bad look, let me tell you. These meetings can be pretty exhausting, because they're typically scheduled back to back all day. Don't be shy about asking for a bathroom break, even if you don't need to go. It can give you a couple minutes to center yourself and also to check over your cheat sheet with notes on your next couple meetings.

Meetings with faculty: These can vary tremendously, and there is no general advice here other than the above. You might want to prepare by looking over a couple of their recent papers. But be prepared for the unexpected. At one of my interviews back in 2009, I remember my first 1-1 meeting was with a faculty member whose office was a dark room in the basement. I walked in, and the first thing he said to me was “Do you know this thing Twitter? It is so stupid. But I can’t stop watching it!” And so we talked about that for a while. More typically, the faculty member will tell you about their research and ask you a few questions about your own. Again, be sure to have your one minute research speech ready in case they didn’t make it to your seminar (or if it’s before your seminar). Look for points of common ground. Or common musical interests. Again, it’s all over the place. Main thing is to just be yourself.

Meetings with students: Often, there will be a lunch with grad students from the program. These meetings can be quite informative. Again, treat people with respect. I have heard of candidates just sitting on their phone the whole time during the lunch, which comes across badly, and yes, the committee will find out. Students can also be a window for you into different aspects of the institution, and also indicate the sort of trainees you would have in your lab.

Dinner: Usually, you will have dinner with a few faculty. Like I said, you are always on, and dinner is no different. The goal is to get to know you in a more informal setting. As always, be yourself, but of course keep it professional. It’s also a good opportunity for you to see how much your colleagues like each other.

Aftermath

Interview done! And now you wait.

The most standard next step is a second interview, which is usually an indication that there is substantial interest in moving forward with an offer. You might discuss your startup and space needs, and at this point, the ball is a lot more in your court. However, until you have an offer letter, nothing is final! I had a second interview that turned out to be very much a second first interview, so you never know. Again, you are always “on”.

Other random thoughts:

- **The process is opaque, live with it.** The process seems opaque, and you typically don’t get feedback if you get rejected. Don’t press on it. You might be able to get some second-hand whispers about the reason for the rejection, but you typically won’t because these reasons are generally best kept secret, for everyone’s sake.
- **Emailing afterwards.** I think it’s not a terrible thing to email faculty you met after the visit, but I don’t look poorly on candidates who don’t send me an email. If you want to follow up on specific points, then feel free to email away! Don’t read anything into the responses to these emails (or lack of response entirely), it doesn’t mean anything. As faculty who gets these emails, I have to say that responding to them is a little awkward, so I tend to avoid it.
- **Don’t sweat little stuff.** I think I heard somewhere on Twitter that speed of reimbursement is a signal of institutional function/dysfunction. Uhhh, that’s silly, if you ask me. Those little things sometimes just are the way they are. Every institution has some administrative ups and downs. I would not judge any place on that.

Negotiating an offer

I honestly have little experience with negotiation, but I did get the following advice: once you have had time to consider, make a list of things you want from the place you want to go to, and say “If you give me these things, I will come”. That can open doors, because the chance of landing a recruit can sharpen the mind, so to speak. It also allows everyone to leave the negotiation with a good feeling. Avoid multiple rounds of back and forth. It can really put people off.

Also, again, don't sweat little stuff. The key is to find an institution that you feel supports you and your goals. The best indicator is their past success at doing so, so see how people have done before you. Startup funds are important, but if it's like 20% less at a place you like more, go for the place you like more. In the long run, that little bit of startup just won't matter.