1. CV advice: How would someone best market their skills gained through a PhD for an industry-related job? Specifically, what is most important to highlight on a CV?
*A: For data science:*

- *Ditch the CV and focus on a 1 page resume. Happy to share mine with anyone who wants an example. Keep experiences concise, remove all or most physics jargon (unless relevant to posting), highlight key tools used where relevant (these tools should tie back to skills you've listed elsewhere), and add numbers to show the impact you had.*
  ***Industry is very impact-focused– they often won't care what you did if it didn't have some flavor of impact***
  *Ex 1: Improved accuracy of subatomic particle identification by X% using OpenCV;*
  *Ex 2: Developed custom simulation software to automate generation of tens of thousands of training images for deep neural network*

- *If you're a fresh PhD, keep your resume to under 1 page*
- *Highlight 5-7 of your biggest achievements per experience*
- *Take a few hours to learn SQL and throw that on your resume. It's a very simple language, and companies are generally looking for folks who know it.*
- *Don't write a cover letter– we don't read them.*
- *Don't put languages or experiences on your resume that you're not able to speak to live. Interviewers may ask you about aspects of your resume during the interview process, so be ready with examples or stories in the STAR format to answer those questions.*
- *Don't include your publications or talks on your resume unless they are relevant to the job posting. They won't generally hurt you, but they also won't move the needle unless they're relevant to the job and create clutter. If you must add them, do it at the very end*

2. Work Life Balance: How much of a balance do you have, between work responsibilities and outside interests compared to your time as a grad student/in academia?
*A: I work 9-4:30/5. Many companies value this balance as part of their cultures. At Wayfair (and many other companies), writing emails or slacks on weekends or after hours is considered bad practice and openly discouraged. So my work life balance is great, but this may vary by company. If you join a 10 person startup, this won't necessarily be true for you.*

3. Getting Started: Was your research group/collaboration supportive of your decision to transition towards industry? Was it easy to find mentors who could specifically help in this transition? Any advice as to how to make those connections?
*A: My group was supportive. I was also lucky enough to have the Insight program to help me transition from academia to DS, but unfortunately that program doesn't exist anymore. I'm happy to act as a resource for anyone looking to transition. Beyond me though, I'd figure out what kind of work you find interesting, what companies are doing that work, and start cold messaging people in those lines of work on LinkedIn. Not everyone will respond, but you'll likely get some traction there. There are also public slack channels you can join to connect with other people in the community.*

4. Skillset: In your experience, do industry jobs care about specific coding/machine learning techniques when they are hiring, or is it sufficient to be able to show that you know at least one area and have the aptitude to problem-solve and learn what is needed specifically for that company once hired?

*A: What you'll be asked really depends on the company, but you'll definitely be expected to demonstrate skills before you're hired. Here are some commonly explored interview topics for data science (will vary by company, but the recruiter will tell you ahead of time what to expect):*

1. *Data structures and Algorithms: Many companies will expect you to solve a data structures/algorithms problem. If you're not familiar with these, I'd recommend "Grokking Algorithms" to get started. Many companies will pull problems straight from Leetcode / Hackerrank easy or medium, so practicing there is also a good idea*
2. *Statistics: Many DS roles will require some knowledge of experimentation and stats. Know all the basics of linear regression (e.g. when is your estimator/standard errors unbiased? How would you design an AB test to assess the impact of some change to our website? What's a pvalue? Ttest? When would you use a non-parametric test? etc)*
3. *Data Fluency: Focused on your knowledge of SQL and pandas. You can practice SQL on Mode. If you don't use pandas, turn some data you work with into a csv and practice common operations.*
4. *ML: This will range from a series of questions about the details of bread and butter ML algorithms (Linear Reg, Logistic Regression, Trees, NN, XGBoost, …), performance metrics (RMSE, precision, recall, AUC, …), business metrics (CTR, CVR, …), etc to case studies. Case studies are generally of the form "the product team comes to you and wants to do X; what do you do?" Ex: we want to recommend 10 products to customers via email; what would you build to do this? What features do you engineer? What metrics do you use to assess your model? If you're looking for practice, go to the company's website and think about what you'd want to know if you were them and trying to make money. Always make sure you understand the business goal before you start trying to answer these questions.*
5. *Behavioral: This will generally be either a resume review or behavioral questions. Be sure to follow the STAR format here. And don't go negative in any of your answers (even if something wasn't your fault)! Always spin things positively, otherwise, people worry the negativity will follow you.*

*A few other parting guidances for searching for roles:*

- *Interviews will generally have a recruiter screen, 1-2 tech screens, and 4-5 "onsite" (virtual now) interviews*
- *If you have a PhD, you should be coming in as a Senior Data Scientist (i.e. not entry level). Titles are not uniform across companies– a "Senior Data scientist" at one company may be a "Lead" or "Staff" position somewhere else. So anchor more on total compensation than title*
- *When I was on the market, the average total comp (salary + bonus + stock) for a Senior Data Scientist in a big city was between 120-160k; you should expect the upper end of that range for larger companies and the lower for startups (FAANGs will be outside this range). Check Glassdoor to gauge salary for the company and role you're considering. If*

*you’re a post doc with a few years of experience managing a lab, you can target much higher total comps, depending on your skill set and experience and the company*