

The job hunt can be summarized into three phases:

1. Getting interviews
2. Performing well in interviews
3. Negotiation

Each of these phases requires some thought to approach and succeed within. I actually think that I over-optimized on phase 2 when the most difficult part for me was actually phase 1.

In addition to my own experiences below, I would also recommend reading the following blog posts / personal experiences:

- [LLM \(ML\) Job Interviews \(Fall 2024\) - Process • Mimansa Jaiswal](#)
- [The AI research job market shit show \(and my experience\)](#)
- [Job Hunt as a PhD in AI / ML / RL: How it Actually Happens](#)
- [Landing your First Machine Learning Job: Startup vs Big Tech vs Academia | Towards Data Science](#)
- [My 2025 Job Search Journey as a Technical HCI PhD | by Yuexi \(Tracy\) Chen | Medium](#)
- [Grace's Interview Prep Doc \(Fall 2025\) | Notion](#)
- [Landing Your ML PhD Dream Job · Scott Freitas](#)
- [LLM Prep | Notion](#)
- <https://space.ong.ac/escaping-flatland> [General advice on job searching, what to think about]

For understanding startup equity, I highly recommend the following post:

<https://www.quarter--mile.com/Startup-Equity-101>

Types of companies

Anthropic, OpenAI, GDM: Struggled to get interviews

- One avenue: "Safety fellowship programs". Extremely competitive
 - <https://alignment.anthropic.com/2025/anthropic-fellows-program-2026/>
 - <https://openai.com/index/introducing-openai-safety-fellowship/>
 - MATS: <https://www.matsprogram.org/>
 - <https://constellation.org/programs/astra>
- Another avenue: Postdoc @ Berkeley, Stanford, MIT

Big tech — Meta, Amazon, Google. Research "SWE"

AI Startups

When to apply:

- Big Tech: start even late fall year before graduation
- AI Startups: They will only interview you if you are graduating within ~3 months

Random Note: I think a good cover letter actually can matter a bit! I heard back more often from places (startups) I wrote a cover letter for.

Obtaining interviews

TLDR: Headhunter / Recruiter > Soft connection + cold email >> referrals > applying online

This was challenging for me. There are generally two ways to go about getting an interview:

1. Applying online
2. Reaching out to someone at the company.

There is also the option of the company / headhunter reaching out to you, which we will discuss later. Generally, applying online yielded poor results, even with “referrals”. I applied to ~100 positions online (see company list in this sheet: [📄 Job prep master sheet](#)). Without referrals, I received interviews from these companies:

1. Patronus AI
2. Quora
3. LightningRod Labs
4. Archer Aviation
5. Anthropic Fellowship Program (not full-time position, kind of an internship)
6. Axiom Math (Probably the coolest company on this list!)
7. Boson AI
8. Letta AI
9. Cartesia AI

Online applications with referrals that **I received interviews** from:

1. Scale AI
2. Got one interview at Apple, applied to 7-10 positions with referral
3. Two Sigma (Quant)

Online applications with referrals that I **did not receive interviews** from:

1. Snapchat
2. ~5 roles @ Google Deepmind
3. Pinterest
4. Mistral
5. Millenium (Quant)
6. ~5 roles Microsoft AI (Superintelligence)
7. Reflection AI (referred to recruiter by USC student)

Notice that there is very little Big Tech on here! When I was looking for a job, hiring seemed a bit slow in big tech. I applied to most of the companies with referrals, but did not really hear back from any interesting avenues. For example, I was only able to get interviews at Meta and Google for “Research Software Engineer (PhD)” roles, which seemed very boring and potentially de-skilling over the next two to three years (these are roles applying ML / working on specific features in actual products like Google Maps, or random things like that). Hardly exciting compared to all the cool startups doing crazy stuff!

Recruiters and Headhunters

Most importantly, change your status on LinkedIn to “looking for jobs”, and update your profile. You want to make sure your profile has important keywords for recruiters who are doing keyword searches. For example, I was looking for post-training roles, so I added a bunch of keywords in my experience section:

=====

- LLM uncertainty quantification and calibration, reasoning models, post-training and reinforcement learning (RL) with Aravind Gollakota and Preetum Nakkiran.
- Post-trained 7B models with verifiers, TRL, vLLM, PEFT for better uncertainty quantification.
- Large scale LLM evaluation pipelines with evalchem.
- SFT and reinforcement-based fine-tuning

=====

Random note: When doing hiring screens, a lot of managers actually asked about my experience running GRPO on a 7B model. I used a single 8xGPU cluster, and it seems that most people do not have this experience when applying to jobs. You can easily rent such a cluster for a couple hours, and get this experience yourself. Probably for ~\$100-200. I would highly recommend doing this if you are interested in post-training roles. Getting wandb set up, looking at loss curves, etc. You can just do a simple post-training setup, and then you should be able to discuss some of the issues you ran into when getting this working, how you did evals, etc. A week or two doing this is probably at least an order of magnitude more useful than another NeurIPS publication!

I also updated my personal website with the following information, although I am not sure how much it helped / if anyone useful views my personal website:

=====

I am graduating May 2026 and am currently looking for full-time industry positions related to

- Trustworthy or human-centered AI
- Uncertainty quantification / calibration for LLMs and other machine-learned models
- LLM post-training
- Ranking and matching algorithms / systems
- Other broadly related topics

Please reach out to me at devic@usc.edu if you know of any opportunities! Thank you so much :)

=====

With an updated linkedin / personal website, recruiters may start reaching out to you. This was how I got interviews at Capital One, Intology, Voleon, and DE Shaw.

However, this is not the end goal. I would say the most valuable connection you can make is what I would call a “super-recruiter” or “talent headhunter”. Here are some example profiles I found on LinkedIn:

- [Callum Williams | LinkedIn](#)
- [Chris Farish | LinkedIn](#)
- [Jessica Bergin | LinkedIn](#)

If you want to work for an AI startup, then your goal should be to get the attention of as many of these recruiters as possible. There is an entire industry of AI startups needing AI/ML talent, and then hiring these headhunters on a contractual basis. The incentives here are very aligned for you as a candidate. In particular, if you join a startup through a headhunter, they will receive a fee from the company that you join. So, they are doing their best to not only get you hired, but also to fight for you to get the compensation you want (since they will only get paid if you join the company!).

There are a couple ways to get in touch with these headhunters. With LinkedIn recruiter, you can request to connect with them directly and then DM them with a personalized note:

“I am a graduating PhD student looking for roles in post-training, LLM evals, and XYZ. I have 10+ publications / preprints in venues such as NeurIPS, ICML, AISTATS, etc. I am looking to join a small, high impact team with great colleagues”,

or something like that. I found that this was not a super high response rate. Instead, I found a little trick. When looking on LinkedIn Jobs, you may occasionally spot weird postings like the following:

The screenshot displays the LinkedIn Jobs interface. On the left, a sidebar titled 'Top job picks for you' lists several job recommendations based on the user's profile and activity. These include roles from Latitude AI, Calico Life Sciences, XPENG, Ludo Robotics, and Google. The main content area on the right shows a detailed view of an 'AI/ML Engineer' job posted by 'Jobs via Dice' in Burbank, CA. The job listing includes details about the role, location, duration, and key responsibilities. The interface also features navigation tabs at the top for various job categories like 'For you', 'AI/ML', 'Top US startups', 'Series A/B', 'Easy Apply', 'Unicorns', 'Gen AI', and 'Remote'.

Top job picks for you
Based on your profile, preferences, and activity like applies, searches, and saves
851 results

- Senior Software Engineer - ML Offboard Models**
Latitude AI
Palo Alto, CA (On-site)
\$179.2K/yr - \$268.8K/yr - 401(k), Vision, +1 benefit
1 company alum works here
Viewed - Promoted
- Machine Learning / Senior Machine Learning Scientist / Engineer**
Calico Life Sciences
South San Francisco, CA (On-site)
\$170K/yr - \$233K/yr
2 company alumni work here
Promoted
- Senior Staff Machine Learning Engineer – Autonomous Driving Foundation Models**
XPENG
Santa Clara, CA (On-site)
\$244.1K/yr - \$413.2K/yr - 4 benefits
5 company alumni work here
Promoted - Be an early applicant
- ML/AI Engineer**
Ludo Robotics
San Francisco, CA (On-site)
\$98K/yr - \$305K/yr
Promoted
- AI/ML Engineer**
Jobs via Dice
Burbank, CA (On-site)
Viewed - Promoted
- Senior Machine Learning Engineer, Performance**
Google
Sunnyvale, CA (On-site)
18 connections work here
Promoted
- Research Scientist Graduate (Multimodal Interaction and World Model) - 2026 Start (PhD)**
ByteDance

Jobs via Dice
AI/ML Engineer
Burbank, CA · 7 hours ago · 7 people clicked apply
Promoted by hirer · Responses managed off LinkedIn

On-site Full-time

Apply Save

PREMIUM
Your AI-powered job assessment

Show match details Tailor my resume Create cover letter

About the job
Dice is the leading career destination for tech experts at every stage of their careers. Our client, Infomerica, Inc. is seeking the following. Apply via Dice today!

Role: Senior AI/ML Engineer
Location: Burbank, CA
Duration: 6 months
Role Overview

We are seeking a highly skilled AI/ML Engineer to design, develop, and deploy advanced AI solutions leveraging LLMs, RAG architectures, and agentic AI systems. The ideal candidate will have deep Python expertise, strong NLP knowledge, and hands-on experience integrating AI models into scalable, production-grade cloud environments

Key Responsibilities

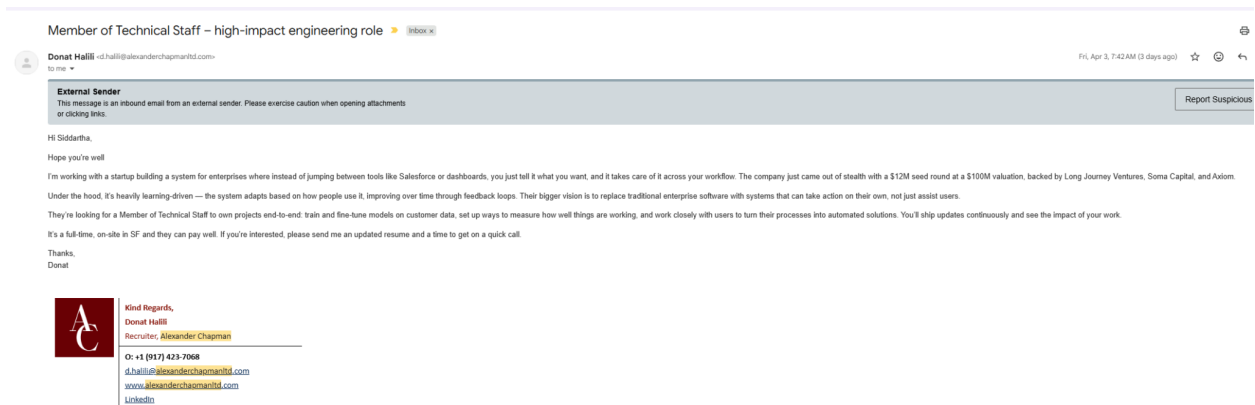
- Design and implement AI/ML solutions using Python and modern ML frameworks
- Develop and optimize Prompt Engineering strategies for LLM-based systems
- Build and deploy Retrieval-Augmented Generation (RAG) pipelines
- Integrate LLMs via APIs (Azure OpenAI preferred) into enterprise applications
- Develop and orchestrate Agentic AI workflows with tool/function calling
- Implement vector search solutions using Vector Databases
- Ensure CI/CD integration and cloud deployment (Azure preferred)
- Establish observability, monitoring, and evaluation frameworks for AI systems
- Collaborate with cross-functional teams to deliver production-ready AI features

You can tell that this is a recruiting / headhunting agency since the company “Dice” is not actually the one looking for an AI/ML engineer, and the company name is called “Jobs via Dice”. With LinkedIn premium, you can apply for up to three jobs a month with the “Top Job Choice”

checkbox. Do this for headhunting job postings, and they will most likely get back to you as long as you state how qualified you are in the note you send along with the job posting.

Once the headhunter gets back to you, you schedule a call with them. They will tell you about one or two opportunities that they are sourcing for. Then, you tell them that you are interested in other companies as well, and ask if they can connect you with more. They will start spamming you with companies that they can set up interviews for you with. This is great! You don't have to accept all the interview requests, but you can do a bit of due diligence, or use these interviews for practice. This was by the far the most effective way to get interviews with 10+ AI startups if you want them!

Once you speak / submit your resume to a headhunting firm, you may also get emails like this:



This looks like spam, but it could be useful! You just want to

Name Brand Matters

Another option that I did briefly consider is to try to get a post-doc at MIT, Stanford, Berkeley, or another name brand school. My friends at these schools pretty much got interviews wherever they applied. It is not unreasonable to consider a post-doc at these schools as a way to get an [@berkeley.edu](mailto:berkeley.edu) email in order to more easily get interviews at e.g., OpenAI. However, nothing is guaranteed, and the opportunity cost is high (1-2 years on academia salary, and if you don't know what you want to research, it may come at a cost to your competitiveness in the market). Ultimately, I believed that I could work on relevant topics in industry to boost my profile if I wanted to go on the job market again in two years or so; furthermore, I wanted to stay in California for personal reasons, and I did not feel like moving across the country.

Personal Network

I have ignored one important avenue for obtaining interviews: personal networking. In Oct-Dec 2025, I would think of random people that I had met in undergrad / PhD / conferences. I compiled a big list of people that I knew, and where they worked. If they worked somewhere cool, I would reach out to them, and state how / where I had met them from.

The more familiar you are with the person, the more they will help you obtain an interview at the company. For example, I had a skiing buddy working for one company. He basically pestered the recruiter every 2-3 days until they gave me an interview. If you have someone inside the company doing this for you, then you stand a much better chance at landing an interview, and once you receive the first interview, you are in the pipeline to completing the entire loop and it's all up to how you perform.

I was able to receive one or two interviews via cold emailing. Here is a cold email I sent that got me an interview at Goodfire:

Hi X!

Hope you are doing well! I am a graduating PhD student from the University of Southern California, and I just applied for an RS role at Goodfire. I wanted to also reach out over email since I know cold applying has a very low shot of my application getting seen.

I really enjoyed reading the work you all did on RLFR (<https://www.goodfire.ai/research/rlfr#>). I worked on similar topics at Apple last summer, trying to post-train models to have better reliability via verbal confidence calibration. I am pretty excited about further work in this direction, and would love to work on similar topics at Goodfire.

I have attached my CV for reference. Thank you so much!

*Best,
Siddartha Devic*

There's also what I call "warm" emailing, which is where you have at least one co-author or point of contact in common with the person you are emailing. For example, I received a response from Alex at Bespoke via the following email:

Hi Alex!

I am a graduating PhD student in machine learning at the University of Southern California (finishing May 2026). I am currently looking for full-time roles in industry, and was wondering if there were any opportunities for research engineers / scientists at Bespoke. I interned at Apple with Aravind Gollakota last summer, and he suggested reaching out to you.

I am broadly interested in working on LLM post-training, trustworthiness, or reliability. My background is generally in foundations of trustworthy AI / ML and theoretical machine learning. At Apple, I worked on uncertainty quantification and post-training in LLMs, which was where I was introduced to the OpenThinker series (and became a fan!). I also interned at Amazon, working on some black-box optimization techniques.

I love the open thoughts dataset, and all the open source work and models that Bespoke has created (I used evalchemy in my last submission as well). I would be excited to contribute to such efforts at the company.

I have attached my CV for more information. Please let me know if you have any further questions about my background. I hope you have a great week!

*Best,
Sid Devic*

Random Note: I told Julian about this strategy a couple years ago, and I did not end up executing it. I believe that it is worth doing for the top companies. Essentially, create a personalized ~1 minute video discussing your background only for the particular person you are cold emailing. Then, send them the video in the cold email. I think it is probably a clear way to stand out against other cold emails. I still think this is a good strategy that someone should attempt! I will keep it in my back pocket for later down the line.

I sent between 20-30 warm / cold emails, with varying levels of response.

Interview Preparation

From mid October through December 2025, I logged all my interview prep preparation in this sheet: [📁 Job prep master sheet](#)

Total number of interview rounds: ~45 across ~22 companies

This is my notes for all my preparation: [📖 Interview Prep Notes](#). I broke down my interview preparation into the following categories.

DSA prep: ~70-80 hours (leetcode list neetcode: [NeetCode 150 - LeetCode](#)). Focus on easy, lower end medium leetcode problems. Don't worry about very tricky hard / harder mediums. In fact, I would limit the amount of time in this category by a lot. Startups generally never asked leetcode questions, only larger tech companies. However, getting good at leetcode can help you become a quicker coder, which is useful when a startup asks you to quickly implement KV caching, or vectorized KMeans, or something like that.

ML Prep: ~15-20 hours ([DeepML](#), [TensorGym](#), asking Claude to generate me boilerplate practice ML interview questions using Numpy, and implement tests to make sure my implementation is correct.) These kind of questions tested your understanding of vectorization, writing efficient and correct numpy code, and basic ML concepts like data standardization, transformers, etc.

LLM Prep: ~22-28 hours ([Stanford LLM course](#), and learning pytorch / LLM training from youtube videos. Also asking Claude to generate questions and tests.) These questions

generally were conceptual (not coding), and asked you to explain key parts of transformer / LLM pipelines, such as positional embeddings,

General LLM / ML Related Coding: ~15 hours (see Claude Coding Practice tab in

[Interview Prep Notes](#))

I continued preparing for interviews as I got them from Jan 2026 onwards, but stopped logging so much because I had the forcing function of the actual interview being scheduled. I would estimate that I spent between 2 and 10 hours prepping for each round of each company that I cared for. This included a lot of asking Claude to ingest the interview description and create for me boilerplate code for practice problems.

By the end, I had a good guess of the kind of problems I would get from just the description that the recruiter sent me. For example, near the end I got this interview description:

The technical phone screen is a 60 minute remote interview in a live programming environment across the course of an hour. You will be asked to solve a technical problem in real time. This exercise is designed to assess a combination of machine learning knowledge, programming skill, and problem solving ability. Programming problems are drawn from the machine learning domain, not LeetCode, but some amount of LeetCode or similar practice may be helpful for comfort with programming interviews and algorithms. Experience with a tensor library like NumPy is often helpful. You may also be asked brief questions about your background and/or ML fundamentals.

Based on this description, I knew that I would get KMeans, Numpy MHA, KNN, standardize data, etc. Indeed, I ended up getting one / multiple of these in the interview, and since I had prepped them the night before, the interview was quick and easy. These are all common numpy questions that you can find on [DeepML](#). I also knew that I should be familiar with broadcasting, since that is usually what they test you on if they state that you may use numpy. To learn tensor broadcasting, I suggest [Tensor Puzzles](#).

Random Note: When getting interview descriptions, if the description says “you may want to be familiar with...” then you should know that like the back of your hand, because they are going to ask you about it!

I also used Claude Code to generate additional coding practice questions in a single repository. Here are the questions I completed (Claude Coding practice tab): [Interview Prep Notes](#) . For the record, I doubt I completed all of these without at least some hints / debugging help. But I try to use that sparingly when interview prepping.

Conservative total prep time: ~100-200 hours. I imagine that I actually spent much more time reading relevant papers, watching tutorials on various pipelines, etc. But I spent around this much time actually sitting down and coding things in order to get my empirical skills up to par.

Here's another random resource on interview prep: [Introduction](#)

Suggestions / Learnings from my Process

Schedule interviews quickly

Multiple times, I scheduled an interview too far forward into the future (~more than 1-2 weeks out). Especially if you do this with the first (usually technical screen) round, you risk the position filling up with someone else. This happened to me with Amazon AGI, Apple, and Cartesia (even though I did well / exceptionally in these interviews according to the recruiter). Essentially, if you hear back from your initial application or referral, you should aim to schedule your interview ASAP since especially for desirable or well-established companies, it is a race to get the position.

Don't bother with quant

I thought doing quant interviews would prepare me for AI companies that I may actually want to work at, but I found that the interview process was too orthogonal to be useful. I didn't need to know any statistics / probability for my main interviews, and quant interviews were still testing leetcode style questions, which have almost completely fallen out of favor with AI startups. Furthermore, general big tech SWE roles asked much simpler leetcode questions, and specific big tech research roles would just ask random LLM / ML implementation questions. So in general, quant firms are not a useful practice for the kind of roles that you may want in AI research (at least this is what my experiences indicated).

Calendars

Make sure you are using a calendar to keep track of interviews, as it will quickly become overwhelming without.

Deciding Where to Work

Don't worry too much about where you are going to work until you have offers. Don't count your eggs before the chickens hatch! Once you have a couple offers on hand that you are seriously considering, I would then ask everyone around you for their opinion. In fact, you should weigh their opinion as almost more important than your own. This is especially true for small startups, which are often beholden to [reality distortion fields](#). I would often interview with a startup that I thought was cool, and then when I talked with one of my friends / co-authors, they would very quickly point out flaws with the company. Beware of founders / CEOs! They are the most convincing people in the world because that is what gets their startup funded!

Here are some questions I asked or wrote down to ask Alex before joining Bespoke.

1. Series A details. Also, what needs to go right in the next 12–18 months for the company to reach the next stage?
2. Who would I mainly be collaborating with day-to-day? How is the team currently structured?
3. What prevents frontier labs from building this internally? Are contracts just a stop-gap until they can figure out how to do these things better by themselves?
4. Could you walk me through how a typical data contract works from start to finish?

5. One of my long-term goals is to eventually lead a team of researchers or engineers, will I be able to do this at Bespoke?
6. What valuation is the equity vested at? What is the strike price of my shares? What is the run-rate of the company? (Ask ChatGPT to understand startup equity as well if you aren't familiar with how it works / it is vested)
7. Publishing / collaborations?
8. What research would most accelerate the company right now / What is the biggest technical or research challenge facing the company right now?
9. What do you see the role of research to be as we grow?

Essays and Readings

Read as many Paul Graham essays as you can, many are remarkably useful in deciding what kind of place you want to work for, and also just general outlook on jobs: [Essays](#).

Moving forward, I really think that the readings from Barath Raghavan's "Crafting a Research Agenda" course will be more important than ever: [Crafting a Research Agenda, CSCI 654, USC, Fall 2023](#). The essays covered include topics like philosophy of science, avoiding cargo cult science, how to think like a PI, "meta-rationality", etc. As we automate more and more parts of our research, thinking about how research should itself be *organized and directed*, will become more important than ever. I believe that these readings give you a good foundation onto which to start thinking about the higher abstractions behind choosing good research directions, and creating valuable research. Furthermore, good writing still seems to be far from being automated, as it requires some clarity of thought and higher level organizational skill which seems to currently be lacking from model outputs.

Apply Broadly

Apply everywhere! You aren't guaranteed an interview anywhere (at least with a USC email lol). Instead, apply broadly, even to places which don't seem name-brand. You may end up finding some smart / interesting people that you really connect with!

Negotiation

Consider using a service like Rora to increase compensation: [Rora: Career Coaching for Elite Technical Talent](#). I did not personally use this service, but if you do not have anyone in industry who can help you negotiate it may be useful.

I had two very close friends mentoring me through the negotiation process, one is an AI researcher in industry, and the other was an ex-VC for Thrive Capital (who raise OpenAI rounds). They helped me craft each email / told me what kind of questions to ask and be clear with my asks. You can see the emails I sent in the Bespoke interview process tab on the left side.

IMPORTANT: If a recruiter asks you what \$\$ you are expecting, never tell them a number. Instead, tell them you are evaluating your market value and need to see what your offers look like. These are basic negotiation tactics!

Also, check out the negotiation tips from Alvin Wan: [Guide to the Job Hunt](#), and use ChatGPT to refine your negotiation emails. Try to do things over email if you can, it may be harder for you to negotiate on the phone with a very pushy recruiter. Also, you can always use the excuse: "I need to discuss this with my partner / family" if you need to get out of a situation of saying yes/no to a certain number.