
Research internships

How to get one, and what to do once you have it

This is written for undergraduates who want research experience and have nobody to ask how it works. It assumes no connections, no famous institution, and no prior publications. Everything here I worked out by getting it wrong first.

1. What a research internship actually is

A research internship You join a laboratory, are given a problem nobody has solved, and work on it under a supervisor. The output is a result: a method that works, a system that runs, a paper. There is no curriculum and no assigned task list.

An industry internship You join a company and are given work that has a known answer. This is valuable, and it is not the same thing. Do not describe one as the other.

The difference that matters In an industry internship, someone knows what success looks like before you start. In research, nobody does. That is the entire skill you are there to acquire.

What it is not

It is not a certificate. It is not two weeks of a workshop. It is not an unpaid position doing data entry with the word “research” in the title. If nobody in the group is trying to answer a question they do not already know the answer to, you are not in a research group.

2. Before you write to anyone

Four things have to be true. Every one of them is about you, not about the lab.

1. **You can build something.** Not a course project you followed a tutorial for. Something that works, that you can explain end to end, and that broke at least once in a way you had to diagnose. A supervisor is deciding whether you will produce anything if left alone for two weeks. Nothing on your CV answers that question except a thing you built.
2. **You can read a paper.** Properly, meaning: you can say what the contribution is, what it assumes, and where it stops. Most people can read the abstract and the figures. Being able to find the assumption doing the heavy lifting is the whole game, and it is what your first email will be built on.
3. **You know roughly what you want to work on.** Not a field, a problem. “I want to do robotics” is not an answer anyone can act on. “I want to work on how a planner behaves when the map is wrong” is.
4. **You have a one-page CV.** One page. Research work first, then projects, then everything else. No photo, no objective statement, no skills bar chart.

If any of these is missing, fix it before writing. It costs you a month. Writing without them costs you the lab.

3. Finding the right group

Where to look

Start from papers, not from universities. Find three papers you actually understood and enjoyed, look at who wrote them, and look at where those authors work now. Read the group's website: recent publications, current students, what they say they work on. If a group has published nothing in two years, it is not active.

How to judge whether it is a fit

The test is whether the group's research programme, taken as a whole, contains the work you want to do, taken as a whole. Not whether one of their papers resembles one of your projects. A one-project match is a coincidence, and it reads like one.

Who to write to

Write to the professor. Not the department, not a generic address, not a PhD student unless you are asking a specific question about their paper. If the group's page says how they want to be contacted, follow it exactly. Some say do not email, apply through the portal. Believe them.

4. The email

Five sentences. Who you are, what you found in their work, what you have built that touches it, a small ask, thank you. Attach the one-page CV. Nothing else.

The sentence that does the work is the second one, and it must contain something specific from their paper: an assumption stated in a remark, a limitation admitted in the discussion, a case called out of scope. This is the only part of the email that cannot be mass-produced, and it is the part that gets you a reply.

What to ask for

Ask for ten minutes, or advice, or an opinion. Not a position. A professor who cannot offer you a position can still say yes to a conversation, and the conversation is where positions come from. Asking for the largest thing first makes the smaller yes unavailable.

One at a time

Write to one group, wait, then write to the next. Thirty identical emails is spam, and academics in the same field talk to each other. If you cannot name what is specific about a group, you are not ready to write to it.

5. After they reply

If they say yes

Ask three questions before you start. What would a good outcome look like in three months. How often will we meet. What should I read first. Most students ask none of these and spend a month guessing.

If they say no

Ask whether they know anyone whose group would fit better. A no from someone who read your email carefully is worth more than silence from ten who did not, and this question converts it into something useful.

If they do not reply

Wait two weeks. Send one short follow-up on the same thread with something new in it: a result, a question, a paper you have since read. Then stop and move on. Silence is almost always a calendar, not a judgement. Professors receive a great many of these emails and answer the ones that arrive on a good day.

6. The first month

Reproduce something before you extend anything

Take the group's published result and reproduce it from their code. You will learn more in that week than in a month of reading, you will find out where the paper is silent about what it actually did, and you will have something concrete to report at your first meeting. If your numbers differ from the published ones, say so. Reporting a discrepancy against your supervisor's own paper is the single fastest way to be taken seriously.

Send progress before you are asked

A short note every week or two: what you tried, what worked, what did not, what you plan next. Keep it factual. Include the things that failed. A supervisor who can see how you think will give you better problems than one who only sees polished results.

Say when you were wrong

If you had a hypothesis and the evidence killed it, write that down explicitly. It costs nothing and it is the clearest signal available that you can be trusted with an open problem.

7. Working remotely

Remote positions are how most people without geographic luck get into good groups. They are harder in exactly one way: nobody sees you working, so the only evidence of you is what you send.

- Reply within a day, even if the reply is "received, working on it, will send Friday".
- Never miss a meeting. If you have nothing, come with the nothing and say what blocked you.
- Put results in a short document, not in the body of an email. Two pages with a figure beats six paragraphs.
- Write down what you agreed at the end of every meeting and send it back. Timezones and accents both cause silent misunderstandings.

8. What to ask for at the end

A letter Ask early, not at the deadline. Ask directly, in a separate message from your work updates. Give them the deadlines, the list of programmes, and a page reminding them what you actually did.

Authorship If you contributed to a result that becomes a paper, ask what your position will be and ask before the paper is written, not after. This is a normal question and professors expect it.

Confirmation in writing A short note on institutional letterhead confirming the dates and the nature of the collaboration. Some universities will not issue one for informal arrangements. Find this out at the start rather than at the end.

9. The failure modes

- **Waiting to be ready.** You will not feel ready. The people who got in were not either.
- **Mass emailing.** It fails, and it burns the group for everyone at your institution who writes after you.
- **Going quiet when stuck.** Two weeks of silence reads as abandonment. Send the problem instead.
- **Only reporting successes.** A supervisor cannot help with work they cannot see, and cannot vouch for judgement they have never observed.
- **Treating a paper as sacred.** Published work has assumptions, gaps, and errors. Finding one is not disrespect, it is the job.
- **Collecting affiliations.** Three names on a CV with nothing produced under any of them is worse than one name with a result. Finish something.

Written by someone who worked this out from a college nobody had heard of. If it helped, or you are stuck on something it does not answer, write to me.