

The Hapless Undergraduates' Guide to Research

By Aaron Gertler

Clinical and Affective Neuroscience Lab, Yale University

Hello!

I don't know who you are, but I'm glad you're reading this!

If you're an undergraduate research assistant, this guide might help you get better at your job and have more fun.

If you're not, I guess you could read it anyway. There are no pictures of baby koalas, but it's at least grade-C procrastination material.

I wrote the first draft of this guide without bothering to Google "guide to doing undergraduate research". That was a bad idea. If you learn one thing from this Guide, it should be to *learn about your work before you start working*.

Anyway. When I turned to Google, I found the fantastic [WebGURU project](#), where Northeastern University researchers have put together a list of guidelines on making your lab work go as well as possible.

If you really want a comprehensive guide to research, read WebGURU. It's great. [Here's that link again.](#)

But I've tried to add some value to the site: I took the general-purpose advice, summarized it, added specific examples from my own research, threw in some jokes, and finished with links to resources that fill in WebGURU's gaps.

The end result?

- 1) If even one person—that's you!—picks it up and saves themselves a few days' worth of frustration, the Guide was worth writing.
- 2) If even one person—also you?—sends me advice I can add to the next edition, it will become even more helpful to future research assistants, like the new kids you'll see hanging around the lab next semester.¹

If you spot anything that seems wrong, or that you think I should add, you can reach me at aaron@gertler.com. *Bon science!*

¹ Undergraduate research assistants last about as long as the average [worker bee](#). The best way to tell these two species apart is that worker bees have stripes.

How to Read This Guide

Start at the beginning and finish at the end.

If the meta-analysis stuff (12-14) doesn't relate to you, skip it.

If you don't like the writing style, [WebGURU](#) really is quite good.

Table of Contents

2-3: Pages you've already read

4-7: Three very important lessons

8-10: Preparing to do research

11-15: Doing the research (warning: mostly meta-analysis)

16-17: Learning on the job

18: References

Disclaimer:

The author of this guide is a Cognitive Science major who never took a physics class and can't remember what "reagents" are. It is mostly about database research and asking good questions as opposed to mixing things, dissecting things, exploding things, etc. For more on those and other topics, read the References.

Disclaimer (2):

The guide is targeted at those who already work in a lab. If you don't work in a lab but would like to start working in a lab, [go here](#) and do several of the things the writer tells you to do. Also, [WebGURU](#) is worth a look.

I guess you could tell your favorite professor that you've read the Guide and will therefore be the World's Best Research Assistant. That might impress them, if they are like [Will Ferrell in Elf](#) and do not know what lies are.

Seriously: if you're looking for a research position, honesty is the best policy. If you're not genuinely interested in a lab's work, $p < .05$ that you will not be completely miserable after the first month.

So You Work in a Lab

Congratulations! Science is a lot of fun.

Whether you're splicing fruit fly genomes, running top quarks through a wind tunnel, or delivering painful electric shocks to subjects who only signed up to get psychology course credit, I think you'll find your work rewarding and stimulating...

...as long as you remember a few things.

Each of these things is important enough to be given its own page, but here's a quick introduction:

- 1) Stand up for your own time: *Making yourself useful*
- 2) Be very careful: *Preventing disaster*
- 3) Lab is for learning: *Having a good time*

Stand up for your own time

Every older person in your lab was once an undergraduate research assistant themselves. They like you and want you to succeed.

However, they also have their own stuff to do.

This means that, for the most part, the task you're working on will be a secondary consideration. It may not have been optimized for speed, precision, etc.

If something about the task seems inefficient, it's likely because no one fixed that part of it before. That's your job! (With help from the rest of your lab, of course.)

You are the person who values your time most highly. Because of this, you are the person best positioned to ensure that your time is used to produce good results.²

² Always ask before you change things, of course. Sometimes, as I work in my dorm room, I'll feel a slight rumbling under my feet and hear screaming from the direction of Science Hill. And I'll say to myself, *looks like someone didn't ask before they changed things*. And then I'll start running in the other direction.

Be very careful

Lab work can be almost hypnotically repetitive, which means that small errors can be committed over and over again until they compound into disaster.

I've seen weeks of work be proven useless—tens of thousands of dollars spent for no results—because a minor error in a computer program corrupted the data in a way that was very difficult to catch.

You likely aren't very experienced in your field, which means you'll have a hard time noticing when things go sour.

But you can minimize the risk of catastrophe by writing things down: the data you're collecting, the steps you take to acquire it, the big questions you're trying to answer, and so on. [Track your progress carefully.](#)

Finally, talk to someone more experienced whenever feel a twinge of uncertainty and can't resolve the question using your notes or a search engine. In the long run, this approach will save you time and help you become a better scientist, while also minimizing the chance that you'll have to read the same paper eight different times.

Lab is for learning

“Learning, learning, learning.”

—[This guide to the benefits of research](#)

Whether you’re earning money or course credit, lab is a job; you should work hard, and accomplish the goals people ask you to accomplish.

But lab can also be the most useful class you ever take.

You’re surrounded by people who have been studying your subject since you were in middle school—or since before you were born. And most of them will be happy to take a break from dry lab work to answer your thoughtful questions.

And even if you work alone, your work is still the work of *science*—you can still pay careful attention to your tasks and write down the things you see.

But we’ll get to that later; for now, let’s move on to the thrilling work of *process improvement!*

So You Have a Research Project

Awesome! The things you help discover will someday be a boon to the human race, and if you're lucky, your name will appear in scientific databases for as long as mankind roams the universe.

If not, you can at least talk about your impressive science credentials in front of people to make them respect you.

But what exactly is your research meant to accomplish? Are you sure?

Know Your Goals

You don't have to get a PhD in order to do undergraduate research.

(That would be a paradox, and science is *never* paradoxical.³)

But you should try to learn as much as possible about the narrow scope of your own research.

This is the part of the Guide where I start talking about *my* research, which could be succinctly and almost impenetrably described as:

Compilation, filtering, and computation for the purpose of creating a meta-analysis on the neurological basis of human craving.

If you had asked me what I meant by that in February 2013, I'd have said:

"Well, um... we're looking at studies where people saw pictures of food and seeing what parts of their brains lit up."

In June 2013, I'd have been more confident:

"When we crave delicious food or addictive drugs, certain 'reward centers' in our brains become very active. We're looking at many papers that specifically compare what happens when people see images of food and drugs to what happens when they see pictures of boring things, like houses and spoons."

³ [Except for when it is.](#)

I knew the work, but not the rationale.

But now, in January 2014, I can answer in ridiculous detail:

“We’re looking for studies that specifically examine what happens in the brains of certain people under one of two conditions: Either they are healthy people looking alternately at pictures of food and pictures of other things, or they are drug users and non-drug users both looking at drugs.

“By looking at the differences between the two conditions in each scenario, and using data from several dozen papers, we can see which brain areas most of the papers found to be very active, and thus provide extremely strong evidence that those areas are partly responsible for craving.

“But to be sure, we first exclude any studies where other things might be affecting the brain, then record many different variables to look for ways that those variables might interfere with our results, then look at the way each researcher measured the data they reported to be sure we trust their methods, then enter the data from each paper in a way that lets us easily analyze it with software.”

Science accomplished!

But still, there’s no reason I shouldn’t have known all that by March. I got wrapped up in the research itself and assumed my methods were sound.

That’s why I’ve read every one of the 60-something papers we finalized between 5 and 15 times.

And why, if I ever meet [Marc-Andre Cornier](#) at a neuroscience party, I will embrace him and thank him for writing such clear and comprehensible papers on what happens in our brains when we see pictures of cheeseburgers. The project really was *far* more troublesome than it had to be.

So how can such repetition be minimized? How can we be sure we’re doing the right thing when we plunge into our painstaking research?

Well, there is one convenient trick...

Talk to the Experts

The best way to make sure your work is useful from the get-go is to know exactly what you're looking for.

And the best way to find that information is to ask lots of questions to your professors, the senior post-docs, and even the veteran research assistants.

Good questions for the head of the lab and experienced researchers include:

- What is the final goal of this project? How will we know when we are done?
- Is there a certain time by which certain parts must be finished?
- Which information is *definitely*, *definitely not*, or *potentially* relevant?
 - Note: If you hear any disagreement between researchers on any of the above questions, resolve the difference *as soon as you can*.
- Was anyone else working on this part of the project before, or on a similar project within the lab? If they aren't in the lab anymore, how can I get in touch with them?

Good questions for people who once did yours or some similar task:

- What problems did you encounter in the course of the research? How did you resolve them?
- Are there any useful tricks you're found for reducing error, increasing speed, or just generally making your life easier?
- Do you have any written notes that I could take a look at?

Once you feel that you know what your work will entail, and that you'll be able to spot some of the most dangerous errors before they jeopardize the science, it's time to start doing the work!

Doing the Work

This is the part of the Guide where every reader is likely to be doing something completely different.

I haven't run any experimental subjects before, I've never slammed particles together at close to the speed of light, and I've never dissected a frog.

What I have done is read a huge number of papers and enter lots of data—the work of creating a meta-analysis (MA), where findings from a diverse collection of labs come together to form the [Megazord](#).⁴

So I'll talk about both of those activities, and recommend that you check the References if this work doesn't describe what you do in your lab.

Reminder: You are unlikely to care about the topics covered on pages 12-14 if you aren't working on a meta-analysis or other database research. Beware.

⁴ No, wait; that's Mighty Morphin' Power Analysis, which is totally different. My bad!

Reading All the Science

Over the course of six months, my lab partner and I reduced 4500 Pubmed results to 800 viable abstracts to 200 candidate papers to the 100 we entered into our database to 55 or so that were used in the final analyses.

This may have been necessary to compile a comprehensive meta-analysis, where every paper counts, but there were strategies we could have used to make the process much simpler.

The same would be true if our assignment had been a literature review, or simply background research on a new experiment. [WebGURU](#) has much to say on the subject, but I'll quickly summarize what I'd do if I were to start on another meta-analysis today:

1) Hire two research assistants in the Philippines

[Outsourcing](#): it's popular because it works. And science is the universal language. *And* they speak English in the Philippines. What could go wrong?

2) Set up a massive spreadsheet

Most of the team's repetitive reading happened because we added new categories to our data table on three separate occasions—each of which also occasioned a new round of double-checking to ensure quality.

I emphasize asking lots of questions to the rest of your lab to save you from the same experience. If you ask ten questions that each take five minutes to answer, and this saves you from spending an hour on one paper you shouldn't be entering (or one subject you shouldn't be running), you save time. The math is on your side.

But even if the lab isn't sure a certain category will be helpful, best to include it just in case. If it turns out you don't need to know the hair color of an experiment's subjects, you just waste a few seconds on each paper. If you have to find hair color later, you wind up checking every paper again to see if hair was even included—and sometimes having to reject those which failed to differentiate between dark brunette and light brunette.

Say it with me: *Lots and lots and lots of categories.*

3) Read other meta-analyses on the same subject

This will help you in a few ways:

- a) You'll get to see which categories past researchers have used, reducing the chance that you'll miss something important.
- b) You'll get a head start on finding papers that match your criteria, which can save you from accidentally excluding papers that serve your cause.
- c) You'll sometimes get to see which papers the authors considered but didn't include in the final analysis. You'll face these decisions yourself, sometimes with the same papers; expert information always helps.

Plus, MAs are just a lot of fun to read: tons of data, useful introductions to certain subfields, and results that sometimes overturn the flashy conclusions of single media-friendly studies. Here's a [cool example with vitamins](#), and an [article that summarizes it](#) (alongside several other MAs).

4) Do the massive literature search

This is the hard part, but it doesn't actually have to be hard. Proper keywords help: My search for "food fmri" caught few useful papers that "food craving fmri" wouldn't have turned up, but gave me much more chaff to filter out (though I also found some cool studies this way—see page 16).

In this stage, you'll hopefully be able to tell whether something is useful from the abstract. Keeping your *definitely (not)* conditions in mind will help. *Definitely not* phrases for our study included "anorexia" and "Prader-Willi syndrome", because both conditions alter brain activation in a way that could interfere with the effects of food craving.

Keep in mind, though, that just because a study isn't relevant to your main analysis, doesn't mean it shouldn't be kept around. After I noticed a few papers discussing the impact of exercise on food craving, I went back and saved them all in a separate section of the main data spreadsheet. We wound up deciding not to analyze them, but for a while we thought we might have enough exercise papers for a small, separate meta-analysis.

At some point, you'll have read all the papers that could possibly be relevant, except for the ones you didn't. (That's alright—science is the art of *good enough*.)

You'll analyze lots of data soon, but the methods will vary widely based on what you're studying. So I won't talk about MATLAB coding here, but I will say...

A Few Words about Data Entry

Not the most exciting task, but someone has to do it.

Music helps. The [Pomodoro Technique](#) helps. On days when your favorite science blog sounds much more appealing than your actual science, [StayFocussed](#) helps.

I recommend entering all the data you can find for a couple of *definitely* papers up front, to get a general idea where everything is and make sure you're not leaving anything out of the spreadsheet.

Then, find the rest of the relevant papers, but enter them in a huge splurge once you're done searching: this will help you enter the hypnotic state where (once you're sure you're doing it right) you can get massive amounts of dull work done without much boredom.⁵

Storing Your Data

If your lab has a standard for this, follow it.

If not, save everything as a PDF, name the files in Author Year format (e.g. Smith 2003), and store them on a Dropbox folder you've shared with the whole lab.

In addition, keep an open sheet in your main database where you place links to papers and tag them with relevant information: whether they're *definitely* or *potentially* papers, this will help you keep track of everything once you know exactly which features of papers justify their inclusion.

(For example, a paper looking at the effects of cheeseburger scent on the brains of obese subjects might get the tags "savory", "scent", and "obese".)

⁵ Once you enter the hypnotic state, there's a small chance you'll start working at 2:00 in the afternoon and stop working after the rest of the humans have [fled the Earth in a giant spaceship](#). But at least you'll have other lost research assistants for company!

When Questions Arise

Ask questions as soon as possible to minimize mistakes. Keep a Google Doc labeled something like “Project Master Doc”. Use it to save the answers to every question you ask. This will be much easier than searching through emails later on, and can help you see the progress of your work (very satisfying!), while also letting you check different senior researchers’ responses against one another.

If a certain essential category is often hard to find, keep the definition in your Master Doc. For us, tricky terms included:

- “ROI” or “region of interest”: parts of the brain that get special attention in some studies
- “Blocked”: a very specific pattern for the stimuli presented to subjects
- “Random-effect” vs. “fixed-effect”: a statistical measure that isn’t often explicitly mentioned in the literature

Nailing down reliable definitions for each term saved us a lot of trouble in the second half of the project, as did figuring out *why* they mattered: once we learned the significance of ROI inclusion, it was much easier to figure out which papers were using it without ever mentioning the term itself.

Which Database Should You Use?

Ask the head of your lab. There are a lot of options out there, especially subject-specific databases that I don’t know well. But if you’re curious, [these papers](#) do a good job comparing PubMed, Google Scholar, and some other popular tools.

In short: Google gets you more results, and PubMed allows for stricter searches, but both are special and lovable in their own way.

Learning on the Job

Your main documents are going to be the spreadsheet where you keep data and the Master Doc where you keep questions.

But I also strongly recommend opening up an “Observations File” where you keep all the cool stuff you discover in the lab.

Notes on lab meetings and research procedure count; so does the time when you’re digging through Pubmed and find a paper on [“maple syrup urine disease”](#), and it takes you a few seconds to get past the sheer terror, Google the syndrome, and realize it refers to smell rather than viscosity.

Science can be bizarre and hilarious, and scientists like yourself get to see that up close. A few choice selections from the file I built this year:

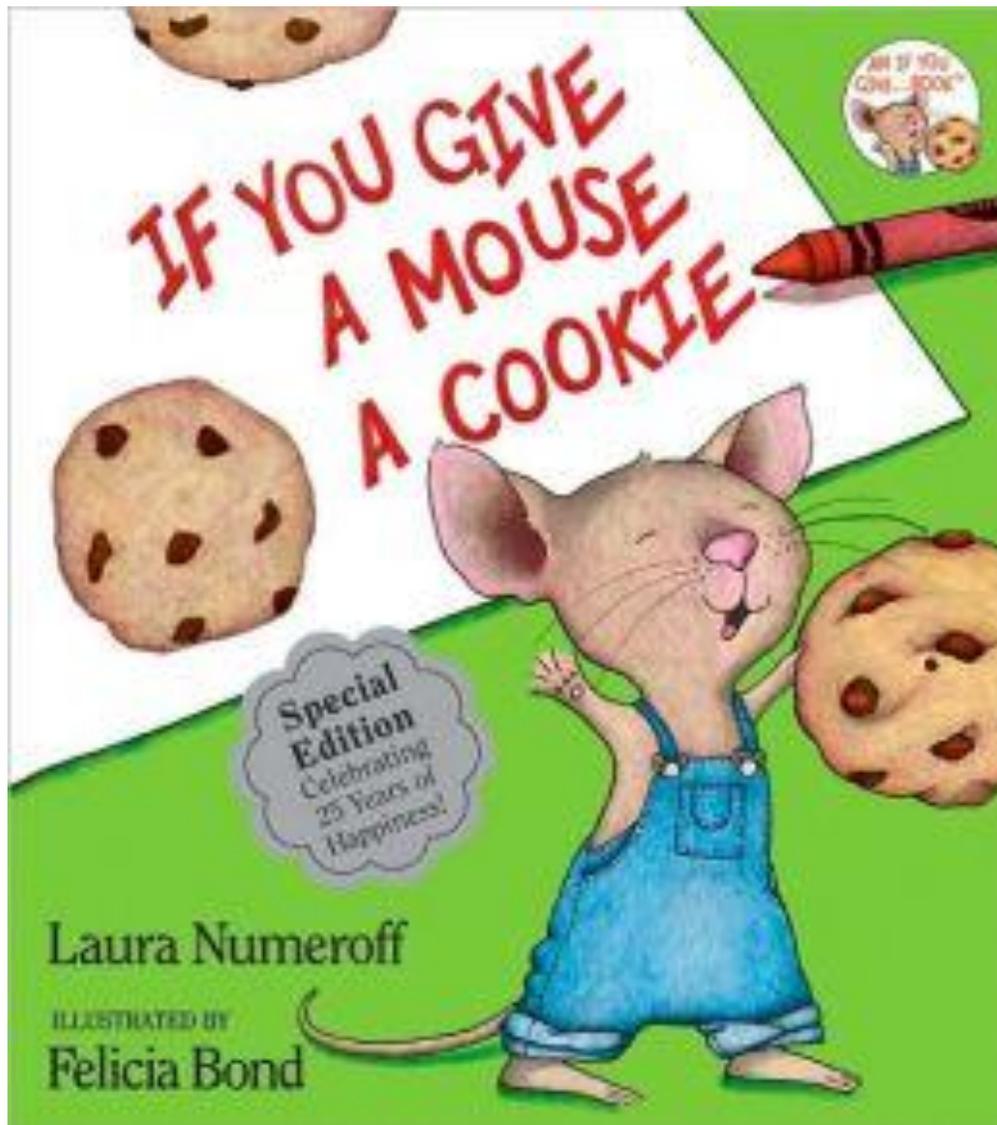
- Someone who smokes one cigarette is more likely to become a habitual smoker than someone who snorts one line is to acquire a cocaine habit.
- Adding pictures of brains to your scientific paper makes people—even other scientists—[more likely to believe what you tell them](#).⁶
- There are multiple common-use databases full of images scientists use for [studies on sexual arousal](#); in other words, researchers at some point had to read vast amounts of pornography just to choose the perfect pictures.
- Just reading about food activates primal centers of craving in your brain—which isn’t so strange, until you remember that, just 4000 years ago, no human being could read any words at all, let alone get hungry *just by looking at markings on a piece of paper*. We learn fast!

In addition to staving off boredom, your Observations File will give you lots of interesting questions for the next lab meeting and help you bond with other people in your lab (everyone loves weird science). Something you notice might even inspire you to start a research project of your own!

And if you start a research project of your own...

⁶ Please do not use this power for evil.

[...you'll need your own undergraduate research assistants to go with it.](#)



Good luck with your research, and check out the references on the next page!

References

WebGURU tackles lots of subjects this guide doesn't cover. Some of the material I found especially useful covered:

- [Frequently asked questions](#) (especially regarding lab notebooks)
- [Independence and self-awareness](#)
- [Moral dilemmas in research](#)
- [Time management](#)
- And, as a bonus, there's an article on [ethical case studies](#), covering some interesting worst-case scenarios (built any unstable skyscrapers lately?)

Independent Research

If you're thinking of writing a thesis, undergraduate or otherwise, [this book](#) seems to be the authoritative source on the subject.

I'm also a fan of [Alan Kazdin's work](#) on research methods in the social sciences—dry, but very thorough. (Disclosure: I've taken a methodology class from Kazdin.)

Hard Science

A [guide to research for chemists](#), but with applications for biologists, neuroscientists, and lots of others.

On [physics research](#), which sounds terrifying and intimidating but also like a lot of high-octane fun! Good luck, physicists.

All the Rest

A solid, simple Top 10 list for [getting started with research](#) (especially “ask for help”, “label EVERYTHING”, and “budget more time than you think you'll need”).

Professor's list of [common undergraduate mistakes](#) (law research, but still useful)

An [in-depth perspective](#) on all the benefits that come with research: “The only way any of us got into science is because someone said ‘yes’ to us.”

[PhD Comics](#): ostensibly for grad students, but just as true to the undergrad experience in some respects.