
General Tips To Drive Your Research At Mila

Victor Schmidt

May 14, 2024

CONTENTS

1	About the guide	3
1.1	Contributors	3
1.2	Contributing	3
2	Introduction	5
3	Meta: external resources to do good research	7
4	At Mila	9
4.1	Doing research	9
4.2	Be a good peer	12
4.3	Help yourself	13
4.4	General First year advice	14



ABOUT THE GUIDE

1.1 Contributors

This guide was put together over the course of several years by (in chronological order):

- Victor Schmidt
- Guillaume Alain
- Tristan Sylvain
- Tegan Maharaj
- Emmanuel Bengio
- Tristan Deleu
- Joey Bose
- Vincent Mai

1.2 Contributing

This is not meant to be a static guide. **You can contribute.**

To do so, just open a PR on the base repository: github.com/vict0rsch/tips-research-mila. You can do so by [forking](#) the repository.

1.2.1 Guidelines

- Try to be succinct and explicit
- Put yourself in the shoes of the very diverse crowd of *new* students: they come from different geographical, social, cultural and educational backgrounds.
- It's ok to have a different opinion from what's already there. If you can explain in a simple sentence why this is the case, it will be more helpful for more junior students.
- The documents are written in Markdown. If you're familiar with Sphinx, know that you'll also be able to use rst.
 - You can use rst-in-markdown with [MyST](#).

Note: Yes this documentation could be `.rst`-only and not mix Markdown and ReStructuredText. But most people aren't actually familiar with the RST syntax so accessibility to new contributors prevailed.

1.2.2 Building locally

The documentation is built and rendered automatically by ReadTheDocs.io.

If you want to contribute and see results for yourself, there's only 3 simple steps:

```
pip install -r requirements.txt
make html
open _build/html/index.html
```

Warning: This is not an official Mila document.

Note: A (bad) pdf version of this document is available [here](#).

INTRODUCTION

This document aims at helping Mila students navigate the research environment and the Mila community. It is a collection of pieces of advice from current and former Mila students and is therefore neither exhaustive nor entirely consistent: people will have different takes on what works or not, how to deal with research, mental health etc. **Read carefully but not religiously: find what works for you!**

META: EXTERNAL RESOURCES TO DO GOOD RESEARCH

- The scientific process
 - Hamming, [You and Your Research](#) June 6th, 1995 (YouTube)
- Scientific project definition
 - John Cleese on [Creativity In Management](#) (YouTube)
- Writing an article
 - Simon Peyton Jones, [How to Write a Great Research Paper](#) (YouTube)
 - [awesome-computer-vision/Resources for students](#) (Github)
- Writing reviews
 - Fredo Durand, 2009, [Peer Review](#) (Slides)
- Writing rebuttals
 - [How we write rebuttals](#), Devi Parikh, Dhruv Batra, Stefan Lee, 2020 (Blog)
- Presenting your work
 - [Giving a Research Talk](#), Fredo Durand (MIT CSAIL) (PDF)
 - [How to give a good talk](#) (Slides)
 - [Designing conference posters](#) (Blog)
- General advice
 - [Lessons from my PhD](#), Austin Z. Henley, 2021 (Blog)
 - [HOWTO: The Reasoning and Learning Lab](#) - Various Artists, 2021 (Blog)
 - [How to avoid machine learning pitfalls: a guide for academic researchers](#) - Michael Lones, 2023 (ArXiv)
 - [An Opinionated Guide to ML Research](#) - John Schulman, 2017 (Blog)
 - [Choose Your Weapon: Survival Strategies for Depressed AI Academics](#) - Togelius & Yannakakis, 2023 (ArXiv)
- Slack
 - [#slow-science](#)
 - [#academic-wisdom-and-advice](#)
 - [#mila-office-hours](#)
 - [#help](#)
 - [#mental-health](#)

- #codingtips
- Experiments
 - Writing Code for NLP Research (Slides)
 - 10 tips for machine learning experiment tracking and reproducibility: Do it yourself approach without additional tooling (Blog)
- Reference managers
 - Zotero
 - Mendeley
 - Papers
 - PaperMemory

4.1 Doing research

- Go to **reading groups**
 - Present, ask all your questions,
 - it is normal that some aspects of a presentation do not appear clear to you: there probably are many others like you, so asking the question will generally be helpful for the group (however, there are also more rare cases where you should refrain from asking more questions, when the speaker is getting too many questions, or you have been monopolising the questioning)
 - Reading groups only work because people contribute to them, both participation and organisation. Everyone appreciates when someone volunteers to present and contributes to the discussion and effort of making the reading group operate well.
- Be aware of your own and other people's **impostor syndrome**: ask your questions, there's a good chance you're not the only one with those. It's how you'll learn, and help others learn too. Think about all those times you were happy someone actually asked the thing you were wondering about.
 - If you are a part of Mila it is because you are **talented**. As such you will be surrounded with other talented people. This is naturally intimidating because people are talented in a myriad of ways, some which you likely are not. The same is true for you! Remember that people look up to you because of the unique ways in which you are talented.
- **Reach out to people**. In the vast majority of cases your peers will be happy to help, discuss your or their research, projects, hobbies.
- **Don't stay stuck**
 - It's normal to face obstacles and uncertainty
 - It's not normal if it lasts for "too long", if it becomes paralysing or if you're starting to get anxiety
 - Identify resources around you: your advisor, postdocs and senior PhD students in your group, or in other groups, Lab Reps, ombudsperson, or even just Slack
 - It is difficult to know if what you're experiencing is normal or not. You need to discuss things openly to have your peers' opinions. Remember: 95% of the people around you have been stuck at some point or have experienced some form of impostor syndrome.
- Do your best to attend **conferences** whenever possible. The research *community* is the lifeblood of science, and you'll find the research community at those conferences. You are privileged to be able to attend conferences, compared with the rest of the planet, so try to make the best out of it.
 - **Meeting** people at conferences and discussing research, e.g. at poster sessions, is more important than attending all the presentations, and conferences can quickly become exhausting.

- Doing some planning before the conference to decide on a few talks and posters you're most interested in can make a world of difference.
- **Classes** are not meant to be magistral demonstrations where interrupting the professor is rude. If you don't understand something during class, there are good chances that 50% of the class doesn't understand either so you're doing everyone a service by giving good feedback to the professor in the form of a well-timed question.
 - If the class is remote, **turn your camera on**. Just like students, professors need to connect with their audience, they are not here to talk in an empty void. Seeing your presence and reactions will help them drive the class and conversations, which will eventually help you.
- **Ego** can be an enemy of scientific progress: we need to keep our mind flexible, be willing to change our opinion in the face of evidence, the arguments we are given. That means we need to truly listen to others and the substance of their arguments. Instead of feeling bad for admitting you were wrong about something, turn that into a quality, be proud of having the mental flexibility to change your mind when reason and coherence demand it.
- A good researcher should be able to play **both the leader's role and the peer/helper** role in a project, with the same grace, enthusiasm and selflessness
- Try to get early **exposure** to important aspects of the research life. Early in your PhD/Masters... try to review for workshops (so as to learn how to do so), collaborate on papers (even if as one of many authors) so as to learn how to write a paper in practice. Try to get opportunities to give small talks (e.g. reading groups), organize events (e.g. workshops): this will help you be more comfortable with similar tasks later on.
- Learn how to say **no**
 - Choose your projects and collaborations wisely, don't get overwhelmed by too many involvements
 - Saying no means making tough choices and focusing your research. Of course it's not just a blunt no: make sure to consider the time you really have
 - Overcommitting is bad for your research, your mental health and others if you end up contributing less than you thought you could
- Learn how to **fail**
 - Be wary of the [sunk cost fallacy](#)
 - Stop when you're not progressing (with caveats, e.g. try to understand *why*)
 - Give yourself time limits
 - Don't give up too quickly, but do give up at some point (finding this balance is where your supervisor(s) and senior colleagues should be consulted)
- It's **okay** if you don't have the next 5 years of your life planned out. It's okay if you don't have the next few months of your life planned out. You can still do research. It always looks like everyone's working on something great. But they have down times too. Talking about it helps.
- When you generate a **plot**, make sure you also save the data that was used for making that plot. You might have to generate a similar plot in 3 months or 3 years, and that jpeg won't help you then.
 - Use PDF plots in Latex
 - Save the script that generated the plot
 - Ideally create a folder for the plot, its script and its data on Overleaf, or on Github
 - When reporting results, add a comment with a link to the wandb/aim/comet/etc. run if you've used an experiment tracking tool (you should)
- Use code versioning (**git**) and store commit hashes, branches and any relevant information to make that result reproducible. For others, but especially for others if you need to produce additional plots or tables in a rebuttal.

- **Save your model checkpoints.** Whatever ends up in a paper should be kept safe (in a project folder, your \$ARCHIVE, your \$HOME, on your personal computer, whatever). Reference the ID of the checkpoint as a comment in the Latex code. Remember \$SCRATCH is not safe, it's purged regularly. This sounds obvious, but here's a scenario you should avoid:
 - Work for months on a paper
 - Save hundreds of checkpoints on \$SCRATCH
 - Write a paper
 - Reference several architectures & ablations & baselines in the paper
 - The paper is submitted
 - The paper is rejected
 - You improve
 - Submit again
 - A reviewer asks for a metric on a model that you had not anticipated
 - It's been 6 months since the original experiments, and because of the **regular purge of files on \$SCRATCH**, some (all?) checkpoints are gone
 - OR: you referenced a dozen models in tables of the first version so you don't know which checkpoint to use
- **Talk about authorship early**
 - Dealing with authorship is one of the most delicate issues in research. People often wait for the paper submission to talk about it, in a rush, under pressure, and it can go wrong. At the very least it can be stressful.
 - One way to circumvent that is to talk openly about authorship ambitions as soon as the team and the project are stabilised and a paper is in sight (even if in 6 months). You can trigger this discussion yourself, you don't have to wait for it to happen.
 - * Example: "I'd like to talk about authorship in the next meeting" so that people have time to think about what they want. And then
 - "I'd like to be first author. I'm open to having co-firsts if anyone also wants to be first. I plan on working a major part of my time on the project and contribute [here] and [there]. I'm saying this for two reasons: 1/ because I care about this project, I will work a lot on it and I will want it to feature it in my PhD thesis and 2/ because I also want you to feel like you can tell me that I'm not delivering first-author work so I can adjust."
 - "I will work on this project 5-10 hours a week so I'm ok being middle author. I want this project to succeed but I cannot/don't want to lead it myself."
 - "As a postdoc, it will be important for me to show I can supervise students and a project altogether. I know PIs are generally last authors, but this project would be particularly important in my future career".
- Graduate school is a time of great **freedom to learn** beyond your own specialty. You are surrounded by experts in many fields, at Mila and in Montreal's universities. If you can, take the opportunities to discover new disciplines and better understand their research problems and their approaches. You may find an opportunity to apply ML to a new problem, or an insight that will shed a new light on your own research problem. In any case, you will develop a better general scientific knowledge.
- A crucial skill you will learn to develop is to identify (and shape) **ideas worth pursuing**. You can do this by drafting them as if writing a paper, including mathematical definitions, assumptions, hypotheses, and potential results. This process will help you uncover gaps in the ideas and give you an estimation of the necessary work, before even starting experiments.

- Create a professional **website** so that people can learn about your work and research interests. As a good starting point, you can create a profile on the Mila website, with your contact information and a short bio. For more flexibility, it is recommended to create your own website, that also includes an updated CV and a list of your papers. You can use a framework such as [Jekyll](#) to create your own static website, and you can host it on [Github Pages](#).

4.2 Be a good peer

- **Spend time at Mila** physically. Good random encounters will happen if you give them a chance to happen — think about it like a Binomial(n, p) where you want n to be large when p is fixed.
- Spend time in the **kitchen**, join and start discussions and ask people open-ended questions
- Be **mindful** and **polite**.
 - Try to notice if someone is being excluded from the conversation or if someone is talking over others, and steer the conversation in a better direction
 - e.g. “I noticed X got interrupted a couple times; X I’d like to hear your perspective!”.
 - If someone doesn’t seem to want to talk, you could say something like “Thanks for the chat, if you’d like to talk more you can find me _____”, and let them choose to continue the conversation if they wish,
 - * or simply ask “Do you want to keep chatting about this?”.
 - If someone asks to be left alone, do absolutely leave them alone. If someone talks to you and you don’t feel like chatting you could say something like “Thanks it sounds interesting but I’m not free to chat right now.”
- In discussions online or in person, always try to give the **benefit of the doubt** and respect others’ opinions and differences. Pause or disengage and look for trusted support if you’re feeling uncomfortable or upset.
- **Organise** your own reading groups or socials.
 - These don’t have to be public events! E.g. grab friends or classmates and read a paper together every other week.
- A lot of research ideas come out during *downtime* (e.g. GANs came to Ian Goodfellow while he was having a beer with friends), provided that you’ve prepared your mind for these ideas to grow (i.e. going to class, reading papers and so on).
- Amazing **internships** generally happen because you connected with a researcher at a conference or in some other social setting.
 - Practise talking about your research and interests in a concise and engaging way, and practice finding interesting connections between your interests and others’ work, but also don’t treat your interactions like a business opportunity. Don’t make it all about yourself, rather about shared interests.
 - Human connection, empathy, and social skills are very important for productive collaborations (internship supervisors are going to work intensively with you; everyone wants that to be an enjoyable experience).
- Try to adopt the mindset that everything good that happens at Mila is because 1) *someone* decided to launch that initiative, 2) *nobody* forced that person to do it so you have to be thankful that they organised the chaos a little bit, 3) that person is your **peer**, and for that reason you are also empowered to start amazing things to make Mila better.
 - In other words, Mila is **yours**. Don’t just sit back, organise something, share thoughts, include people around you, show up.
- Different students might have **different experiences** due to their supervisors and affiliations and the particular research group they find themselves in.

- You can learn a lot by watching how one of your colleagues **works**. Same goes for reading someone else's code.
- Mila is great, but there are a lot of processes at Mila that are broken. Some of the dysfunctional systems only manage to get the job done because we assume that Mila members are **smart** and can compensate for the errors or irregularities. Be that person, be smart and contribute to fixing what you see to be broken. Be patient and kind with other people working in the sometimes-dysfunctional system.
- When collaborating with colleagues, *test your code* (within reason) before sharing it with others. Nobody likes wasting an afternoon because someone else committed bugs in the shared repository. Specify requirements.
- Ask your friends, desk neighbours or people you just had a coffee with to **explain** what they work on. Take 30 minutes on a whiteboard to learn what others are doing. And present your work to others to get feedback.
- Don't ask to ask, just ask -> <https://dontasktoask.com/>

4.3 Help yourself

- **Ask for help.** When asking for help, be respectful of the other person by making it as easy as possible for them to help you.
- You can generally get to meet any person by taking the time to **read** (some of) their papers and asking them questions that show that you've done work (i.e. you're invested in them, in some way).
- Find a way to work in a **sustainable** way. Pay attention to your mental health, sleep rhythm and listen to your body.
- You have a certain amount of **freedom** at Mila that regular employees at office jobs don't have. Make that arrangement work for you. Keep in mind that the easiest way to get help and direction from your supervisor is to work on the topics that they're interested in the most.
- **Montreal** is a city where you're generally allowed to be weird as long as you're respectful to people.
- Sometimes, work at Mila may feel like it becomes an overwhelming part of your life. It is ok to feel this way, but remember that your life is not limited to Mila, and you should always strive for a healthy **work / life balance**.
- If you're sick, stay home and **rest**. You're not going to impress people by showing off about how much you can neglect your health.
- Focus on your research in a sustainable way: avoid **burnout** by resting and understanding that good research takes time, but also try to eventually close your different projects (e.g. by submitting to a workshop/conference, or giving a talk on your partial findings).
- Deadlines (conferences/workshops) are a good way to structure your research throughout the year, and helps with giving the extra push to finish a project. However, be aware that it is ok to not submit something to that big conference that everyone submits to, and to wait for your project to mature and be the best work that can be: there is always the next conference. As an alternative, there are venues that are accepting submissions all year long (journals, or TMLR).
- Understand that research can be taxing on your **mental health**. Take some time to acknowledge the pressures you feel (peer pressure, pressure to get results, to publish...), and that feeling that way is both very common, but also can be addressed by talking to others, relaxing.
- Remember that you are (quite likely) at Mila to **learn**, and be a student. Producing knowledge is an important part of research, but **learning** to produce knowledge is a *more* important part of your graduate experience. When in doubt, look back at how much you've learned in the last N+1 months/years.
- Take **vacation** time. There is no good default, but 4 weeks / year is low. People regularly take 6-8 weeks / year. 1 day off does not count, it takes time for your brain to disconnect. Conferences are not time off.
 - Again, be a good peer: warn your collaborators in advance about when you'll be off.

- Sometimes, it's easier to be off when others are off too: there won't be emails or Slack messages that will trigger FOMO (fear of missing out) or stress. This will typically be mid July -> mid August and around Christmas. And there will be less work to catch up with when you come back.
- Know thyself! Embrace your way of working. There is no one size fits all recipe to grow as a researcher. Research and life would be boring otherwise. Accept there might be ups and downs. Don't give up (within decent limits of your well-being).
- It's okay to consider quitting your program. It's important to give yourself the option, so as not to feel trapped. If you choose to do something else, it's perfectly fine! If you decide to stay, great! It would be for the right reasons.

4.4 General First year advice

- Your “first year” is not necessarily 12 months, it can be a bit more, the following are pieces of advice for the beginning of your PhD, typically before your predoc.
- Research is a very self directed endeavour. This means you have to set your own goals, own agenda, and own execution plan. Be your biggest critic but at the same time be your biggest champion.
- Your first project will not be revolutionary and you might not be as excited about it as solving a grand AI challenge. That's ok. The important thing is the first project should challenge you technically and help you develop research skills (coding, writing, mathematical maturity). More importantly it should help develop your own research taste. Finally, it should lead to a publication or a paper submission (rule of thumb: after a year of you starting).
- A good rule of thumb in your first year. Have 1 project that you are leading—i.e. first author, and one project that you are supporting (co-author). Ideally the project you are supporting is led by an experienced researcher who could be a senior Phd student in your group or at Mila, and whom you can learn from.
- Your first year is a great opportunity to try a bunch of different research areas and fail spectacularly. Failing is normal and expected; not trying research is not. Do not fall into the trap of only doing courses and feeling that you do not know enough to start a project.
- The main goal in your first year is to develop a research taste. This means finding something that you have an unhealthy research obsession with and more importantly one that you have the skillset (or can develop the skillset) to tackle. This is the key to doing great work. See Paul Graham's essay on this topic: <http://paulgraham.com/greatwork.html>
- Research becomes exponentially more fun when you work with an able collaborator. **Find someone that you can work with.** This often involves being able to ask them stupid questions without feeling like you are wasting their time. Additionally, this could mean both of you are ready to get your hands dirty on technical details—often without any concrete goal outside of getting a deeper understanding of the problem.
- **Talk about PhD / MSc expectations with your supervisor** after a year. *Even if* you already had that conversation when they hired you, or at the very beginning of your studies: it's ok to have that conversation again. The words and their implications will make much more sense now that you're more experienced.
 - The 3-paper rule for a PhD is a rumour. But rumours often originate from some form of reality: check with your supervisor what *they* expect from you.
 - * Have that conversation with senior PhD students in your group: they may phrase your supervisor's expectation differently...
 - Take notes, or have the conversation *via* email so you can refer to it in the future, even just for yourself in a year or two
 - *In general* (do have that conversation anyway!) to graduate with a PhD you will have to show not only research output but also *leadership*. So it's not about “that many papers” but rather about your supervisor feeling confident they can take you in front of a jury who will agree that you have a proven track record of leading good research projects.