

This document summarizes considerations for students looking to apply¹ for graduate research opportunities at Montréal HCI and should be read before meeting Professor Arawjo (me). I have written it to help students who understandably may not know what human-computer interaction (HCI) is, and may also have not clarified their own goals for pursuing a graduate degree. I ask that any prospective student considering our graduate or PhD program read this document. The following sections provide some questions to ask of yourself, some clarifications about the present research interests of Montréal HCI group, expectations of PhD students, and broader information that may factor into your decision-making process. Note that not all of this info may be relevant to your circumstance.

Why a graduate degree?

A graduate program can be very rewarding, in that it offers you the freedom to deeply engage in a topic, become an expert, and build interfaces that help people outside of the constraints of capitalism. (You may even develop tools that are released to the public that help real people!) A PhD can especially be a springboard to rewarding employment in research positions in industry or academia. However, it is important to be aware that a graduate program is not for everyone and comes with its own limitations, chiefly the upfront financial impact and the challenge of being self-sufficient when pursuing a research topic. The ultimate goal of a PhD in particular, as I see it, is to expand one's mind and build one's independence in responsibly pushing the boundaries of human knowledge and thinking critically, without much outside support. One way to look at the trajectory of a PhD student is to go from nodding one's head at everything a professor says, to questioning the professor, to outright disagreeing and being unafraid to voice their opinion firmly, respectfully, and with well-reasoned arguments.

Why do *you* wish to pursue a PhD/Masters degree?

Although it may be too early to think long-term, a PhD is a long-term commitment (Masters less so, but usually more of a financial investment). You should ask yourself what you foresee after the degree —are you aiming for an industry research position? An academic job? Do you just want to learn and explore and mess around for a few years (which is fine!)? Be honest with yourself. A good litmus test is to ask yourself, *could I pursue the project I want to in my off-time? Do I need a graduate study for it?* If you sense you need deeper expertise or understanding, or relish such learning, then, a graduate degree may be right for you. If you rather sense you could pursue it, but something *else* is holding you back, then you should consider otherwise. We do not have a lot of time on this Earth, so use your time wisely.

Why do you wish to pursue graduate research *in HCI*?

HCI differs from other computer science disciplines in its focus on *people* and *design*. Human-computer interaction encompasses a wide range of methods, from qualitative research in ethnography and interview studies, to building technical innovations in VR. However, always the focus is on humans (“users”) and how they use or adopt computer interfaces (whether those interfaces are built by us or by others). The Design research in HCI can challenge what students from science and engineering backgrounds consider to be “research”: in short, *we consider design to *be* research in HCI*, whereas other CS fields consider research to be technical

¹ If you are **already** a PhD/graduate student and wish to collaborate, this document is not as relevant, but see the section above about HCI and the “Current and Future Research Interests” section below.

innovations largely divorced from human factors (design for these fields is an implementation detail, if it is considered at all). That is not to say one is better than the other, but that each discipline has different values, methods, and goals.

The current and future research interests of the Montréal HCI Group

At its core, Montréal HCI is about *interfaces to help people in the world*. A healthy amount of our work is about *designing and building* interfaces. There are also opportunities for *theoretical and empirical research contributions* (e.g., see the pre-print on *Antagonistic AI* by Cai et al.), but usually in service of a longer-term goal to build an interface to help people.

Right now we have two+ research foci:

- **Programming interfaces:** Interfaces to help people program computers, powered by AI technology. In particular, we are interested in interfaces which involve *handwriting/sketching* and/or *natural dialogue* over a design, programming, or software engineering task. For instance, a past application was handwriting quantum circuits inside Jupyter Notebooks. A current project delves into how AI can help game developers ideate (design) a game in the early stages, and develop a game design doc.
- **AI model auditing interfaces:** This follows from ChainForge (chainforge.ai), an open-source tool for helping people audit LLM behavior and prompt engineer. In general, we are interested in building powerful tools to help people —developers, potentially even lay persons —audit the behavior of LLMs. There are applications to education, software eng., etc.

There is also a budding interest in *AI for game design and development*, which is related to the first research area and highly relevant to Montréal, one of the top cities for games (and game studies) in the world.

External Fellowships and Funding

If you have a **fellowship** (here, external funding guaranteed by a government or organization), then these research areas can be relaxed to cover any topic under the HCI purview. Examples are: a project on accessibility research, a project on embedding sensors in a forest that applies ML, or a project on social media in the Global South. These are important areas, but know that Professor Arawjo may not be able to fund them outright. Holding a fellowship could also open you up for opportunities to shop around for a program and advisor that is right for you, and often pay higher than normal PhD stipends. Note that most [fellowships in Canada](#) must be applied to *after* one is already entered a PhD program, i.e. during the first or second year. Look around and see what you could apply for. UdeM has many internal sources of support as well.

Expectations of a PhD student

The expectation of a **PhD student** at Montréal HCI is to have published at least three first-authored papers in HCI venues by the time one graduates: at least 2 conferences, and 1 journal and/or conference. HCI is a conference-led discipline, with arguably the most prestigious venues being CHI and UIST, alongside the mixed journal/conference venue of CSCW. If we publish three papers together, you have a very healthy chance of obtaining a PhD. Although this seems like a lot, I have found that it is better to put expectations upfront, in writing, so that students can gauge their progression in a PhD program and have solid targets to aim for.

Working with Industry

Professor Arawjo has a healthy number of connections with industry researchers. Should we work well together and produce research papers, there is a good chance you could obtain a summer internship or part-time internship down the line. The benefits of internships are connecting with other researchers, building expertise, and also supplementing your finances during a PhD. The best internship is one that pays well and also allows you to conduct research relevant to your PhD. For Masters students, Professor Arawjo also knows many people in the field of HCI with active PhD programs and could make recommendations.

The Differences Between Canadian and U.S. PhD Programs and Funding Models

It is important for applicants to also understand the differences between U.S. and Canadian institutions with regards to funding, and the pressures this can put on students and professors.

U.S. and Canadian PhD opportunities, on the whole, differ in three ways:

1. The first is that **Canadian PhDs (including at UdeM) usually require a Masters degree**, although exceptions can be made if the student is otherwise exceptional. U.S. PhDs, by contrast, tend to recruit students fresh out of undergrad, and a Masters can be earned along the way to the PhD (or one can leave early with a Masters). In general, Canadian PhDs also have a significantly reduced courseload, with the vast majority of time spent on research.
2. The second is that **Canadian PhDs take on average less time than U.S.** (for reference, a PhD at Cornell Information Science takes on average 6 years). If you are diligent in your studies, it is possible to complete the PhD in four years.
3. The third is **the funding model for PhD students**, which interacts with the first two differences. See below for more details.

The way the Canadian funding model for PhDs works is that the professor guarantees a certain amount from their funds (at least \$15000 a year in Montréal, which has a relatively low cost of living) that is topped up by the department through TAships or other internal funding sources or grants. This differs from some U.S. institutions, such as Cornell and NYU, that fund a PhD through an amount provided by the university² which is detached from a specific advisor's grant funding. This "guaranteed" amount of U.S. universities is actually a requirement that the PhD student holds a TA (teaching assistantship) for each semester where they have no graduate research funding. TAships can eat into time, lasting 15 hours a week or more. Nevertheless, the U.S. funding model can offer greater freedom to the student to pursue their own interests than a specific grant or professor may provide. And, at Canadian universities, depending on the funding of the professor, one may still be asked to TA.

The difference arises from the fact that, in Canada, a large portion of, if not all government grant funds can be *used directly* by the professor —there is little "overhead," due to the Research Support Fund. By contrast, at U.S. universities the actual grant amount a professor obtains is reduced by the overhead of the university —and the overhead can be quite large, easily 50% or 60% of the money that the student actually sees. So, even though NSF funds appear larger than equivalent Canadian grants, a large portion of an NSF grant is taken off the top. Effectively

² This is not exactly true; see the next paragraph.

then, U.S. professors are actually putting in extra funds a year for students, even students they do not pay as GRAs; however, the funding is indirect and subsidized by other professors' ability to obtain grants. Operating costs may also be why hyper-competitive U.S. institutions require professors to secure a pre-specified amount in grant funding (say \$1 million) before they can obtain tenure. A good rule of thumb is that Canadian grants and fellowships are, on average, easier to obtain than the hyper-competitiveness of their U.S. counterparts, but the amount one receives may on average be less (a risk/reward trade-off).

What Matters? A Good School or a Good Advisor?

One last thing to mention is how you should structure your search for PhD programs.

What matters in a PhD program are four things: that *you feel supported and happy* (by peers, professors, the broader community where you live, family, etc.), that *you and your advisor get along*, that *you do research that matters to you*, and that *this research has impact* (through publications, advances in technology, contributions to public debate, etc). Your goal in looking for a PhD program is to find an advisor, institution, and location best set up to support you.

This may be a more controversial topic, but I see some prospective students have a hyper-focus on getting into the “top” schools for their PhD studies. I believe this drive comes from a mindset first cultivated in secondary school and by society at large, that is arguably true for undergraduate studies but increasingly less true for PhD studies. I have seen students even willing to give up freedom and overlook red flags to get into a “top” PhD program.

The truth is that, at the end of the day, no one cares if you come from a “top” institution if you haven't done good research, compared to a person who has high-profile publications/impact. Us professors know this intimately. As a person who has experienced both a good undergraduate institution but later worked at the so-called “top” schools, I have seen that the apparent prestige of a university, from the outside, has almost no bearing at all on the quality of your advisor and your experience. At best it is an incredibly fuzzy signal. The majority of value that prestige provides comes from *outsider perceptions* of that prestige: that is, it is easier to make network connections. It also seems that, sadly, more prestigious schools can cultivate advisors and situations that are problematic—whether the advisor having so little time that students feel isolated and alone, or outright abusive advisors that cultivate hostile, demeaning environments. That is not to say this is a hard-and-fast rule or that other institutions are somehow free of problems—and certainly, as I mentioned, there are upsides to prestige, and you'll find good people everywhere—but rather to warn you that you should consider all factors when making decisions about PhD programs. A good strategy is to reach out to other PhD students and prior collaborators of a potential advisor. How was their experience? Try to gain a holistic picture.

Conclusion

If you've read this far, congratulations! :) You have the healthy reading appetite of a graduate student. If after reading this document, you wish to chat with Professor Arawjo about research opportunities, please reach out to his email at [first name].[last name}@umontreal.ca (all lowercase). Include a copy of your CV, and briefly describe your background and interests.

Written March 28, 2024. This document may be updated in the future.