

Inspired by similar documents by Chara Podimata, Sean Sinclair, mentorship plans at UMich (thanks !) edited with chatgpt (I can be wordy !)

My advising style:

The tools that I work with/develop are:

- Statistical machine learning
- Causal inference
- Sequential decision making under uncertainty (MDPs, ...)
- Optimization (under uncertainty, contextual, stochastic, robust)
- Modeling / operations management
- econometric mindset (credibility revolution, RCTs, ...)

The span of all of this is “data-driven decision-making”. You don’t have to be interested in *all* of these, but you’ll need to be familiar with at least one for us to work together effectively. We may be an especially good fit if you are interested in connections among them.

What types of questions am I interested in currently?

Methodological focus:

- structure in offline/online reinforcement learning
<https://arxiv.org/pdf/2406.08697>
- Optimal adaptive data annotation for robust LLM-informed causal inference
<https://arxiv.org/pdf/2502.10605>
- Learning and decision-making about individuals over time

Areas of interest re: social impact

- program evaluation & algorithmic accountability
[BIRS workshop report](#)
- service delivery in the public sector
- interdisciplinary collaborations for human/ai collaboration in algorithmic decision support

This is just to give you a flavor of types of questions I am interested in: while sharing at least one tool in common is a hard constraint, we will work together to find questions of mutual interest.

Projects: The ultimate goal is for you to become an independent researcher.

One of the best ways of getting better at doing research is by doing research. In the beginning, we will talk about shared interests and work together to find something tractable/interesting that will also allow you to get hands-on experience with conducting and writing research and important tools and techniques (methodological, theoretical, and applied).

However, as your PhD progresses, I am there to support you in developing ownership of your research interests and contributions. Though along the way, I may suggest concrete project

ideas related to your interests, having you become an independent researcher is the end goal.

Importantly, our time is limited, so we will have to triage many good project ideas in pursuit of great ideas. When choosing between projects, I somewhat prefer depth over breadth; especially fruitful areas for you to work on. Collaboration is encouraged, however it is best for you to have a strongly developed thesis topic that is clearly yours.

Operations Program-specific advising: Qualifying exams in the first year cover optimization and stochastics, requiring some self-study beyond coursework (do this with your cohort!). The first summer offers a chance for research collaborations, and it's beneficial to connect with potential advisors beforehand. The summer of the first year is the research portion of the qualification exam, and is a good opportunity to work together on a project to experience research with a particular advisor/topic area. Given that, it's usually good to spend some time prior to summer matching/connecting with potential advisors. Our program encourages you to get started discussing research right away (balanced realistically with your courseload).

Finally, it's important to note that in the Operations program, you are absolutely free to work with multiple faculty (so long as they are available/interested).

Meetings: We'll meet one-on-one once a week, which is a good place for us to work together on projects and also talk about broader career questions, goal-setting, etc. I'm also available for quick questions on Slack/email. I think it's crucial to develop your own sense of research taste, so we'll also discuss what you find interesting and why; and how that meshes with what the field finds interesting and why. Especially if we are working remotely (mostly for external collaborations), I might even suggest twice-a-week shorter meetings, though the default is weekly. We may meet more frequently before deadlines as needed to revise papers.

Group: Once there is a critical mass, I plan to set up group meetings to facilitate community-building among students I am specifically working with. Group meetings will comprise of either reading a paper together, or presenting ongoing work or other types of practice presentations. The frequency and time of group meetings will vary.

In the meantime, if you are interested in causal inference, we are co-organizing a reading group on causal inference (including students across DSO, CS, Econ). If you are interested in social impact, I highly recommend getting involved with Mechanism Design for Social Good: I have been in the working group on Inequality. I've met great friends and collaborators through MD4SG.

Schedule: My working hours are 9am-5pm, but I may work evenings near deadlines (usually writing/editing). I don't expect off-hour responses and try to reply within a day. I need a few days to review materials.

Deadlines: There are plenty of deadlines in academia. While I may suggest venues and potential deadlines for completing a project, you are free to suggest other high-quality outlets as

well. Meeting a deadline is a commitment from all of us. I expect that if we have agreed to complete a project for a deadline, we will both work together to do so with a growth mindset.

Inclusion: Diversity, equity, and inclusion are incredibly important to me as a faculty member. Recently in our fields, really fantastic affinity groups have grown, which I encourage you to check out if relevant. (Women in OR/MS or Machine Learning, Minority Issues Forum; Black in AI, Queer in AI, LatinX in AI, Muslims in AI) I am always learning and open to feedback as to how I can act as a more inclusive advisor.

Receiving feedback: I seek your feedback on how advising is going for you and what does/doesn't work.

Giving feedback: Feedback is crucial for growth. I will distinguish between higher-level considerations vs. nitty-gritty tiny things that I'll show you to do research effectively. Please feel empowered to request certain formats of receiving feedback.

Authorship/authorship order: Authorship in OM/OR typically follows alphabetical order. For interdisciplinary work, we'll discuss authorship expectations early.

Publication venues: We target top OM/OR journals and major conferences in statistical machine learning (Neurips/ICML/AISTATS) and algorithmic fairness (FaCCT/EAAMO). Since MS/OR accept extended versions of conference papers, I personally find it useful to send initial versions to conferences and extended versions to journals, to expand visibility of results, put deadlines to partial progress, expand the use and impact of research, and finally to develop polished, complete projects.

I commit to fund conference attendance for accepted papers you co-author with me. Our program funds your attendance to ~ one conference a year.

My expectations for students:

Honesty and integrity: I expect that you conduct yourself honestly and with the highest standards of academic integrity. There is never any need to "cover anything up", we have to get into the weeds on the work, I'm on your team, mistakes happen all the time: I want to help you fix any bugs.

Learning the field: I encourage you to identify and pursue projects in areas where you have some inherent curiosity.

Meaningful/targeted effort: Good research does not happen without work. However, "work smarter, not harder": there are endless ways to put in effort that does not lead to outcomes. We will prioritize and triage during meetings. I do expect you to work towards some kind of progress between meetings (thinking about the problem, background research, reading papers/learning

new things, partial progress, trying things out, simulations), though this depends on the stage of the project.

Work-life balance: a PhD is a marathon, not a sprint, and I also expect you to take care of yourself and enjoy life outside of the PhD as well. While I will suggest deadlines to keep us both on track in making progress, take these as suggestions: you are free to push back, and they serve primarily to keep us both on track.

I strongly encourage you to take vacation (~4 weeks a year). I do encourage disconnecting from email/the internet/cell service, please just let me know in advance of your travel/disconnecting, so we can ensure upcoming deadlines are uninterrupted. Great times include around the winter and summer: academia also slows down around these breaks.

Kindness/collegiality: We're in it together, and I will strive to be kind. (I do tend towards being upfront so as not to waste your time). It is best when we can work together openly.

Respectful conduct: I expect all participants of the group to conduct respectful discussions. In addition, your conduct in external settings reflects on us as a group. Although it is easy for academic discussions of other work to become critical quickly, I do expect you to express yourself openly and honestly, but always respectfully and generous towards others. That doesn't mean avoiding sharp insights, but expressing them respectfully and graciously. This is actually an incredible academic superpower that can help you establish and maintain your reputation.

FAQ

Why our program in the Marshall School of Business?

- The culture of our group is fairly technically-minded (but still with an emphasis on impact!). Our department houses an operations group (which I am in) as well as a statistics group. Given the growing importance of data-driven decision-making, this is a unique advantage!

Advantages:

- **Full funding throughout your PhD** (with a modest teaching requirement of two semesters, to support your preparation for potential academic careers)
- Flexible opportunities to work with multiple advisors
- A healthy academic job market with few grant-writing obligations and generous compensation
- Strong engineering programs across the university and opportunities to build collaborations
- Note that given trends in recent years of engineering PhD candidates going on the business academia job market, much of student research is similar

More broadly, the academic community in operations management/operations research greatly values principled design and analysis of theory/methodology, but also relevance to practice and modeling aspects. This appreciation of *both* real-world context *and*

technically innovative/interesting work is somewhat rare in other technical communities.

You might be concerned about a smaller student cohort. My thoughts are: while you can make up for this by being entrepreneurial and getting to know students in relevant other programs across campus (ISYE, Computer Science, Econ), it's much more difficult for you to make up the difference in funding compared to an engineering program.

Theory or applications?

I work broadly from theoretically-guaranteed methods to applications, so we can find a good balance for you. It's best to develop theoretical foundations for developing principled methodology during your PhD: these are golden years for focusing and digging deep. But having impact in important areas is important as well.

The most important thing I've learned is that it can be difficult to do both in the same project, so keeping a "research portfolio" mindset of striking the balance you seek across your overall suite of projects. During our advising, we would talk about building a portfolio that suits you.

How can I learn about interdisciplinary topics like causal inference and machine learning?

I'm delighted to be in the Data Sciences and Operations department with multiple causal inferencers! There are fantastic other resources and departments across the university as well in machine learning theory and econometrics.

We'll start by reading some foundational texts/lecture notes on areas of interest, if courses are not offered in the exact years when you are focusing on courses.