TEACHING & MENTORING STATEMENT

Varun Gupta

Computer Science Department
Northwestern University.

mail@varungupta.info
http://www.varungupta.info

Teaching Philosophy I view teaching as a tangible way to create immediate impact, as opposed to the longer time horizon for research. My principle philosophy is simple: focus on the needs of the students (how will the material be useful to them), pay attention to the context (their backgrounds), and deliver the material in the most effective manner to bridge where they are now to where I want them to be. The first two (needs and context) vary in each setting so I address these for MBA vs. PhD teaching separately below. However my method of delivery has a common theme: be as visual as possible, use toy examples, and ask questions with hints to allow the students to discover the answers themselves. From following these rules over the years, I have come to realize that very often less can be more. It is more important that students leave the room engaged with the topic, and with the self-confidence that they can master the material by themselves, than trying to fit a predetermined amount of material in the available time. Another principle I follow is to be respectful to each student: I put significant time in learning the names and backgrounds of my students within the first 2-3 lectures and I am extremely open to questions during lectures. In fact, one of the consistent positive feedback in my MBA teaching evaluations has been my openness to questions during lectures and answering each question with courtesy. My hope is that this individual attention conveys to the students the fact that I believe in their ability to succeed and as a result they push themselves a little bit harder.

Teaching Experience

MBA Teaching I have been teaching *Managerial Decision Modeling* course at the Booth School of Business since 2012. The broad aims of the course are to teach students the basics of optimization modeling (linear programming, integer programming) and decision making under uncertainty (decision trees, Monte Carlo simulation), and build optimization and simulation models within spreadsheets. Over the many years of teaching this course I have tinkered a lot with the material, shedding a large portion so as to cover the most useful concepts in great detail, and adding topical cases such as decision making under adversarial environments to prepare them for applications like cybersecurity and hedging climate risk which are emerging as areas employers are focusing on. A change I would still like to make is to introduce an end-of-quarter capstone case competition. This would give the students a chance to experience the process of going from a real world problem to an optimization model, and better appreciate the power and limitations of the tools they learn during the course. In Fall 2024, I am also experimenting with an AI Teaching Assistant to help students with basic questions on the material covered and Excel troubleshooting (which can be accessed at `https://www.varungupta.info/tAlly-o.html`) and has received positive feedback from students. In Spring 2024, I taught the core *Operations Management* course to the MBAi cohort at the Kellogg School of Management. The content

as well as delivery of this course was in quite a bit of contrast to the Decision Modeling course, and found myself enjoying the engaging discussions that the material engendered.

PhD Teaching I taught *Stochastic Performance Modeling* during Winter 2014 quarter. The course was also cross-listed as a Computer Science course to encourage participation from the students. The course was initially intended to be a methodology course with the aim of teaching students the basic stochastic/queueing models for service systems. In 2014 the Statistics department stopped offering a course on Stochastic Processes which was a required course for our Ph.D. students and acts as a prerequisite for a Stochastic Modeling course. In light of this development, I designed the course carefully to instead teach the basic tools in analysis of stochastic processes with queueing models providing an application testbed where all these tools could be brought to bear. In hindsight, I felt that this was perhaps a more fruitful approach to stochastic process theory. By applying different tools to the same application domain, students learn to view the same problem from different lenses (as Markov chains, as sample paths, as stochastic recursive sequences), each lens with its own power and limitations.

Starting in 2018, I stopped teaching the Stochastic Performance Modeling PhD course since I felt our PhD students were not really pursuing stochastic modeling as a research area. I myself was changing my research focus around the same time from queueing theory to design and analysis of online algorithms more broadly. It was a perfect opportunity to design a new course where I could expose students to the different perspectives that operations research and theoretical computer science communities take on online decision making under uncertainty – from Markov Decision Processes, to competitive analysis under adversarial assumptions, to online bandit optimization, prediction with expert advice, secretary models, prophet inequality, information relaxation methods, and (distributionally) robust optimization. I have been teaching this course since then, and find that teaching this course and my research reinforce each other quite well.

Teaching interests

Given my primary research interests are in probability theory, stochastic optimization, algorithm design, and mechanism design, the courses I would most enjoy teaching are also along the same lines. At the undergraduate level and masters level, I would be particularly interested in teaching courses on **probability and statistics for decision making**, **service operations**, **introduction to stochastic models/queueing theory**, **stochastic dynamic programming**, **game theory and mechanism design**, **randomized algorithms**, and **simulation modeling**. At the PhD level, I would like to teach or co-teach courses on **stochastic processes**, **mathematical toolkit for operations research**, and **online learning and control**.

Mentoring of students

My goal in mentoring PhD students is to prepare them to be independent researchers who have their own research vision, who have a high bar for their work, and who are effective communicators of ideas. During my first interactions I tend to present three of four ideas that I may have been thinking about pursuing at the time and some relevant papers on each idea. Sometimes I may present a tentative model I have in mind, but usually I like the students to be involved in formulating the model itself since

the students have more ownership of the problem in this case and hence are more engaged and excited about research. In the end it is them who will have to stand up in front of an audience to present and defend their work. No two students are the same – some students are technically strong and are happy to view research as solving a series of puzzles. For such students I formalize an initial model and then we work closely on designing algorithms or proving some structural results. Some students are more motivated by the application context, in which case the initial meetings are focused more on finding a tractable model where I tend to use my intuition to filter out some ideas. For the puzzle solvers, the hope is that they gradually develop a taste for picking problems through the process of research (I was like this). For the application oriented students, I find that they naturally pick up the required tools with very little guidance because the problem itself is a natural motivator.

However I do not believe mentorship requires co-authorship – the duty of a mentor is to facilitate the student's transition into a productive researcher. To this end, I remind my students to not let the published results set the bar for what they consider acceptable research, but instead to be ambitious and develop their own metric of what good research is. I practice the same, and hope the students learn from my example. I also strongly encourage them to write at least one solo-authored paper by the end of their PhDs so that they are forced to introspect about their own research interests and philosophy. During this endeavor, I give them feedback on their ideas and writing, so they know that I am supporting them and not letting them hang out to dry or that I do not want to work with them. I encourage my students to talk to and collaborate with other students since this would be their ecosystem of collaborators when they start their lives as junior faculty. I encourage them to attend conferences even if they are not presenting any results of their own to get exposed to the current trends. Above all, I follow an open door policy – my students can come and find me anytime they have ideas to bounce off, or feel worried or concerned about research or any other events in their lives. I am very sensitive to the fact that power dynamics can distort channels of communication between an advisor and advisee, therefore I make an effort to use humor and be self-critical in front of my students so that they can speak candidly. I can not be an effective mentor if the students do not believe that I have their best interests in mind.

Ph.D. students co-advised/advised

Lisa Hillas (graduated 2023) - University of Auckland
Zuguang Gao (graduated 2023) - University of California Irvine
Deniz Akturk (graduated 2023) - University of Washington, St. Louis
Ming Yu (graduated 2020) - Citadel
Luyi Yang (graduated 2017) - Haas School of Business, U. California Berkeley

Masters students co-advised/advised

Yuwei Luo - PhD student at Stanford GSB
Nipun Thakurele - PhD student at NYU Stern School of Business