I am an experienced instructor and research supervisor. I have taught Introduction to Machine Learning (undergraduate and graduate) and Introduction to Reinforcement Learning (graduate) courses at the Department of Computer Science, University of Toronto (check CV and website). I have trained 25+ HQP at different education levels, including one graduated PhD who is a faculty at the Oxford University and six current PhD students, two of which will be graduating by the end of 2024.[1]

Teaching Philosophy

My teaching goal is to assist students in developing a solid foundation in an area of knowledge, in particular machine learning (ML) and reinforcement learning (RL), and to guide them to acquire effective learning skills in the process. Their learned knowledge and skills should allow them to creatively contribute to a diverse set of challenging problems. My research trainees should become independent researchers who are at the top of their field and are willing to contribute back to the research community. This is challenging though! Students have different goals, backgrounds, and styles of learning. A single fixed recipe will not serve them all. To achieve my teaching goal, despite these challenges, I follow the principles of *growth mindset* and *deliberate practice*.

The **growth mindset** postulates that a person's talents and capabilities are amenable to change and can be developed through learning and practice [Dweck, 2006]. For example, while one might not initially be good at theoretical machine learning, they can become an expert after a few years of practice. **Deliberate practice** involves working towards well-defined goals aimed at pushing one beyond their comfort zone. Its success requires regular and immediate feedback by an expert [Ericsson and Pool, 2016], in this case, myself. I emphasize to my students that failure is not only an option, but also an essential part of the learning process. This approach is, of course, natural for someone familiar with RL.

Additionally, I aim to be a compassionate instructor and supervisor who provides an inclusive and low-stress environment for my students. I realize that people face challenges in their life and have ups and downs in their performance. Moreover, I believe that an effective teacher should be an effective learner too. As an RL researcher, I constantly experiment with different teaching styles, deliverables and evaluation methods, as well as how I lead the Adaptive Agents Lab (Adage) and supervise my students.

Teaching Experience

My **Introduction to Machine Learning** course provides a solid and comprehensive foundation to ML. The graduate-level offering of this course attracts many students from both CS and non-CS departments, including electrical, mechanical, industrial, civil, and biomedical engineerings, and mathematics — an evidence of my experience in teaching to a broad spectrum of students. I have created a YouTube Channel for public dissemination of the material.

I developed a completely new lecture-based graduate-level **Introduction to Reinforcement Learning** course. This is the first time such a course has been offered at the U of T. What differentiates it from many other RL courses around the world is that it gently builds a solid mathematical foundation of RL for students who don't have any prior RL background. As I did not find any of the existing textbooks satisfactory, I wrote a book called Lecture Notes on Reinforcement Learning (166 pages), which I am in the process of expanding to a textbook called **Foundations of Reinforcement Learning**. All lectures are publicly available on YouTube.

I am currently teaching the undergraduate course **Neural Networks and Deep Learning** in Winter/Spring 2024. Its content is gradually posted within the next two-three months.

I have the principles of growth mindset and deliberate practice in mind when creating the content, designing the deliverables, and lecturing and interacting with the students. By including a variety of homework assignments, in-class quizzes, research project, paper readings, take-home test, etc., I give students many opportunities to engrain their knowledge and get feedback. Table 1 summarizes the scores

[1] This document covers my teaching philosophy and experience as of early 2024.

Table 1: Summary of course evaluations with mean (median) format. The score map: 1: Not at all; 2: Somewhat; 3: Moderately; 4: Mostly; 5: A great deal.

Question	**Courses**				
	ML-U (2018)	**ML-U (2020)**	**RL-G (2021)**	**ML-G (2021)**	**ML-G (2022)**
I found the course intellectually stimulating.	4.0 (4)	4.5 (5)	4.5 (5)	4.1 (4)	4.2 (4)
The course provided me with a deeper understanding of the subject matter.	4.0 (4)	4.6 (5)	4.5 (5)	4.3 (4)	4.5 (5)
The instructor created an atmosphere that was conducive to my learning.	3.7 (4)	4.3 (5)	4.5 (5)	4.0 (4)	4.4 (5)
Course projects, assignments, tests and/or exams provided opportunity for me to demonstrate an understanding of the course material.	3.8 (4)	4.4 (5)	4.2 (4.5)	4.2 (4)	4.3 (4)
Course projects, assignments, tests, and/or exams improved my understanding of the course material.	3.6 (4)	4.4 (5)	4.3 (5)	4.2 (4)	4.2 (4)
Institutional Composite Mean	3.8	4.4	4.4	4.2	4.3

I have received for each of my courses, indicating the effectiveness of my teaching. I have publicly released the student evaluations on my website.

Future Courses: I am interested in developing new courses on topics such as *Model-based RL*, *Statistical Learning Theory*, and *RL Theory*. I view developing a new course as an opportunity to fill my knowledge gaps in a subarea of knowledge and to distill that knowledge in a most pedagogically clear and coherent way.

Supervisory Experience

I have worked with a wide variety of students with different interests, skills, and ambitions. My goal, however, is always to help them become independent researchers. A new research student needs to learn that the recipe for success in research is different from success in a course. They often feel desperate and question their abilities and talents. The growth mindset helps them to understand that they can improve their research abilities. The deliberate practice, guided by my supervision, helps them push beyond what they can do, again and again.

When a new graduate student joins my lab, I explain my research program and share my notes on a research topic with them. My notes are fertile ground for a trainee to learn and grow. By the time they understand the notes, they have already gained much insight into the problem and are ready to push the idea forward by implementing, modifying, and coming up with related novel ideas. Gradually, they need less of my help in initiating their research projects. My help would then take the form of guiding them identify their most promising ideas and providing feedback on their progress.

The Adage lab is an environment that allows my students to openly share ideas, get feedback, practice their communication skills, and start new collaborations. We have weekly 1h one-on-one and 2h group meetings and regularly scheduled reading groups. This is in addition to the RL Seminar series that I founded and the RL Reading Group that I co-founded at the Vector Institute.

References

Carol S. Dweck. *Mindset: The new psychology of success*. Random House, 2006. 1

Anders Ericsson and Robert Pool. *Peak: Secrets from the new science of expertise*. Random House, 2016. 1