bio/nrod08
theoretical neuroscience

Course Syllabus
Winter 2019

instructor

Dr. Blake Aaron Richards – blake.richards@utoronto.ca – Office: SW525A

teaching assistant

Jordan Guerguiev – jordan.guerguiev@mail.utoronto.ca – Office: SW525

prerequisites

(1) NROC34H3 or NROC64H3 or NROC69H3
(2) MATA29H3 or MATA30H3 or MATA31H3
(3) PSYB07H3 or STAB22H3
Please come and speak to me if you do not have the required prerequisites.

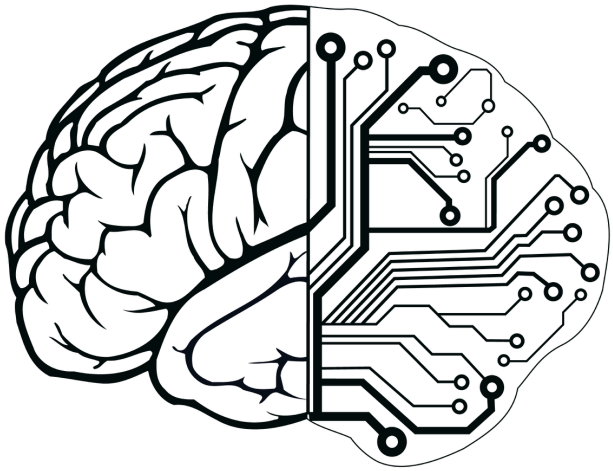

overview and objectives

The brain is a sophisticated information processing device. In fact, despite all our advances in computer science, it is still the most powerful general-purpose computer on Earth. The computations that occur in this incredible, energy efficient computer contained in our skulls underlie our perceptions, thoughts, decisions and memories. Theoretically, if we could understand the operations of this computer, then we could read and write thoughts, create artificial minds, and use brain machine interfaces to extend our minds with technology.

The field of theoretical (or computational) neuroscience is devoted to understanding the brain as a computational device. Using mathematical models of neurons and circuits, theoretical neuroscientists seek a formal understanding of the brain that will allow us to manipulate it in the same ways that we can manipulate other computers. In this class, we will take a multidisciplinary approach to theoretical neuroscience. In addition to learning current mathematical theories of neural computation, we will tie the theories we learn back to biology by studying modern papers that seek a unification of biology and theory. By the end of this course, you should have the beginnings of a formal understanding of how the brain processes information and learns from experience. The specific learning objectives for the course are:

1. Students will understand what computation is, what an algorithm is, and how they could be implemented by the brain.
2. Students will become familiar with neural network models and the theory of parallel distributed processing.
3. Students will understand how information and probability are linked, and be familiar with the fundamental concepts from information theory and probability theory.
4. Students will be able to apply probability theory and information theory to interpret neuroscience data and think about neuroscience experimental design.
5. Students will understand the relationship between learning and optimization, and be able to identify at least two learning algorithms for neural networks.
6. Students will be able to distinguish between supervised learning, reinforcement learning and unsupervised learning, and explain what each can be used for.
7. Students will understand how neural networks can store memories and make inferences about the nature of human memory based on neural network models.
8. Students will be able to write simple programs in Matlab.
9. Students will be able to read, critically analyze, and discuss scientific papers from the area of computational and systems neuroscience.

These objectives will provide the foundation for any future work in academia or industry related to neural information processing, brain machine interfaces or artificial intelligence. There are many jobs in these areas at this point in history, and a lot of active research. Now is the right time to gain an education in theoretical neuroscience.

But, doing well in this course will require keeping up with the material and pushing yourself out of your comfort zone, especially if it has been a long time since you have done much math or statistics. A few important tips:

- You may not understand everything you read in every paper, but it is important that you understand the basic concepts. Discuss the readings with other students in the class. Discussion will help you to understand things you struggle with on your own.
- If you are struggling at all with the math and stats, come to office hours, and be sure to attend every tutorial. The tutorials will provide you with the chance to do a few basic math problems to help you get comfortable with the core concepts, like probability distributions, vectors, non-linear functions, etc. We will build on these concepts throughout the class, so it is important for you to understand them.

- Get started on the programming assignments as soon as they are released, especially if you are new to programming. They will be designed to be do-able for everyone, but if you are new to code, it will not be possible to do well if you procrastinate. You will need to give yourself time to read computer code and understand it.

asking questions

Please ask questions! As many as possible! However, to make both my and the TA's life easier, please ask questions in lecture, at office hours, or in tutorial. Alternatively, post your questions to Quercus so everyone can see the answer and we do not have to repeat ourselves. We will generally not answer questions sent to us by email.

lectures

Fridays from 9-11AM in SW316

The first hour of the lectures will be a standard lecture format (though questions during lecture are welcomed). In the second hour, we will have a group of students present an academic paper to the class (see the assignments below), followed by a class discussion led by the presenting group. **Note:** *there is a grade for participation in the paper discussions, so even if you are not presenting the paper it is critical that you read the paper and engage in the discussion.*

tutorials

Mondays from 2-3PM in SW316

The tutorials are there to get you up to speed on algebra, calculus, statistics and programming. They will also be where you will learn about your programming assignments (see below) and work on them with the TA. **The tutorials are not mandatory but they are critical.** A student with *a lot* of experience in math, statistics, and programming might be able to skip the occasional tutorial, but everyone else should attend every tutorial. Moreover, everyone should attend the tutorials where the programming assignments will be discussed (see tutorial schedule below). Note that there will be two assignments that will involve writing computer code in Matlab, and the classroom is equipped with laptops with Matlab on them. You can also get the student version of Matlab for $10 from the university here: http://sites.utoronto.ca/ic/software/. If you do not know Matlab, do not worry, part of what the tutorials are for is teaching you the basic Matlab knowledge you need for this class.

office hours

Mondays from 10:00-11:30AM in SW525A (back room)

Feel free to come to office hours to discuss class material, get help with the assignments or just to discuss theoretical neuroscience.

course requirements and evaluation

The final grade is based on three assignments, a midterm test, a final exam, and a mark for participation in the post-lecture discussions. Here is the specific break-down of the marking scheme:

Paper presentation assignment	15%	Date: NA, variable (see below)
Programming assignment 1	15%	Date: **due on February 26th**
Midterm test	20%	Date: **February 15th**
Programming assignment 2	15%	Date: **due on April 2nd**
Final exam	25%	Date: TBD
Participation in discussion	10%	Date: NA

The midterm will take place in-class, during the second hour of class. The final will be held during the standard exam period.

assignments

One of the assignments will be a group assignment that will involve reading a paper and making a 30-minute presentation to the class about it (see more below). The other two assignments will be programming assignments where you will have to complete simulations of neural networks in Matlab. Details of the two programming assignments and initial computer code will be posted on the 3rd and 8th weeks of class.

paper presentation assignment

In this assignment, groups of 2-3 students will be required to prepare a 30-minute presentation on one of the assigned readings. The readings are available on Quercus. We will have one group present each week, in-line with the reading dates listed below in the lecture outline. The presentation must be done using slides (e.g. PowerPoint) and answer the four following questions:

Q1) Background: *What is the gap in scientific knowledge that the paper is attempting to address?*

Q2) Results: *What data, model and/or analyses does the paper present, and how does it help us to fill the gap in our understanding?*

Q3) Limitations: *Does the paper completely fill the gap you identified, or does it leave it incomplete? If it is incomplete, how so?*

Q4) Future directions: *What new questions does the paper raise and what should other researchers do to address them?*

After the presentation, the group will also be responsible for leading a class discussion on the paper. You can do this however you see fit, but one approach would be to ask the class specific questions about the paper's implications and interpretations. The assignment will be marked out of 15 marks, with the following breakdown:

- 2 marks for success in addressing Q1 above
- 2 marks for success in addressing Q2 above
- 2 marks for success in addressing Q3 above
- 2 marks for success in addressing Q4 above
- 4 marks for clarity of the presentation
- 3 marks for leading a successful discussion

I will leave it up to you to pick your groups, and the paper you want to present. However, note that **you must decide your group and paper by January 15th and email me your decision**. Anyone without a group and paper by January 15th will have it assigned to them. **Note**: one of the groups will be required to present on January 18th. However, that is an easier paper.

programming assignments

As stated above, there will be two programming assignments, the first of which will be released in week 3, the second of which will be released in week 8. For both assignments, you will receive a backbone of Matlab code that you will then fill in. You will also be required to answer a few questions about the code, to demonstrate your understanding. The code and assignment details will be posted to Quercus. Your write-ups and code should be submitted on Quercus by the due date (due dates are listed above in the marking scheme).

lecture outline

theme	week	date	subject	paper
information processing systems	1	Jan. 11	Marr's levels of analysis	NA - Admin
	2	Jan. 18	Neural network models	Jonas & Körding (2017)
	3	Jan. 25	Information theory	Oteiza et al. (2017)
	4	Feb. 1	Bayesian inference	Panzeri et al. (2001)
	5	Feb. 8	Compression and prediction	Körding & Wolpert (2004)
	6	Feb. 15	Summary	NA – Midterm test
learning and memory	7	Mar. 1	Learning as optimization	NA – Review of midterm
	8	Mar. 8	Deep learning	Rumelhart et al. (1986)
	9	Mar. 15	Reinforcement learning	Kaligh-Razavi & Kriegeskorte (2014)
	10	Mar. 22	Unsupervised learning	Schultz et al. (1997)
	11	Mar. 29	Auto-associative memory	Rao & Ballard (1999)
	12	Apr. 5	Summary	Wills et al. (2005)

tutorial outline

theme	week	date	subject
assignment 1	2	Jan. 14	Vectors and matrices
	3	Jan. 21	Probability theory
	4	Jan. 28	Introduction to Matlab
	5	Feb. 4	Assignment 1 and Matlab continued
	6	Feb. 11	Assignment 1 and Matlab continued
assignment 2	7	Feb. 25	Derivatives
	8	Mar. 4	Gradients
	9	Mar. 11	Assignment 2 and Matlab continued
	10	Mar. 18	Assignment 2 and Matlab continued
	11	Mar. 25	Assignment 2 and Matlab continued
exam review	12	Apr. 1	Review session

accessibility

Students with diverse learning styles and needs are welcome in this course. In particular, if you have a disability/health consideration that may require accommodations, please feel free to approach me and/or the AccessAbility Services as soon as possible. AccessAbility Services staff (located in Rm SW302, Science Wing) are available by appointment to assess specific needs, provide referrals and arrange appropriate accommodations 416-287-7560 or email ability@utsc.utoronto.ca. The sooner you let me know your needs the quicker we can assist you in achieving your learning goals in this course.

absence policy

If you are prevented from attending or completing a course requirement (test, assignment, etc.), due to any illness or other circumstance of a grave nature, contact me by email within two days of the missed requirement, clearly stating the reason. In addition to this email notification, you need to complete the 'Verification of illness or injury' form or other official documentation of the grave circumstance and deliver it to Jennifer Campbell in the Biological Sciences office (SW421D) during regular working hours Monday-Friday. These documents will be used to determine eligibility to recover any lost marks. The 'verification of illness or injury' form can be found at http://www.illnessverification.utoronto.ca/document/Verification of Student Illness (VOI) - Oct 27 2016.pdf. Acquaint yourself with its content such that in case of an emergency you can obtain the essential information required, even in the absence of the official form. When you are well, work hard and engage! If you are sick, take care of yourself, do what you can from home and get back into the swing of the course once you are well. Feel free to contact me for strategizing on getting caught up if you get sick.

academic integrity

The University treats cases of cheating and plagiarism very seriously. The University of Toronto's Code of Behaviour on Academic Matters (http://www.governingcouncil.utoronto.ca/policies/behaveac.htm) outlines the behaviours that constitute academic dishonesty and the processes for addressing academic offences. Potential offences in papers and assignments include using someone else's ideas or words without appropriate acknowledgement, submitting your own work in more than one course without the permission of the instructor, making up sources or facts, obtaining or providing unauthorized assistance on any assignment. On tests and exams cheating includes using or possessing unauthorized aids, looking at someone else's answers during an exam or test, misrepresenting your identity, or falsifying or altering any documentation required by

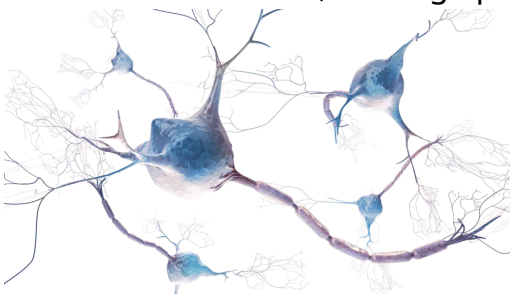

the University, including (but not limited to) doctor's notes. Please avoid academic dishonesty, have confidence in your own ability to learn and grow academically by doing your own thinking and writing! Here's a simple tip: if you ever find yourself wondering whether you are breaching the Code of behavior rules, you probably are. Apply common sense, and make your education the ultimately goal here.