Send To Printer | Close Window

Changelog

Version 1.0 - August 10, 2020

Version 1.1 - August 24, 2020 - added final exam date (Nov. 19) to syllabus

Front matter

Meeting times: Thursday, 3:00 – 4:15pm

Meeting location: Genome Sciences 1377 or virtual

Instructor: Dr. Tyson Hedrick, thedrick@bio.unc.edu, https://biomech.web.unc.edu/

Office hours – by appointment

Course overview

This course explores how an emerging artificial intelligence technology, Deep Learning, is contributing to the analysis and acquisition of biological data. The course is a special seminar with a 1 hour a week meeting time, it has never been taught before and it may never be taught again – such is the nature of special seminars. This means that I don't have a large reservoir of prepared material and the class is in large part a collaborative effort between the students and faculty; while I will lead some early classes I'll also be relying on each of you to help explore this topic area and teach your fellow students.

Course tools and reading

BIOL 680 does not have a textbook, but we will be reading a variety of different sources from primary literature to literature reviews to software documentation and tutorials. We will also occasionally be attempting to demonstrate or test different Deep Learning approaches, and you'll likely want one or both of Python and MATLAB installed on your laptop and/or become a practiced user of Google Colab. However, this is not a programming class and it is assumed that you're already familiar with one or both of these languages. Furthermore, Deep Learning applications commonly benefit from running on a graphics processing unit (GPU), or at the very least somewhere other than the laptop CPU you're trying to use for the rest of your UNC courses. Unfortunately, no typical laptops have the standalone nVidia GPUs that are most useful for Deep Learning. The Google Colab environment does offer limited free access to nVidia GPU or Google TPU (tensor processing unit) instances so we may make use of that resource and/or local UNC resources that would be more suitable for ongoing lab research work. Specifically, consider signing up for an account on UNC's Longleaf computer cluster if you don't already have one – see https://its.unc.edu/research-computing/longleaf-cluster/ Use my name and ONYEN (Tyson Hedrick; thedrick) as the faculty sponsor, bash as the shell, and put "Deep Learning development and testing" or similar as a description of your work. Note that basic Longleaf accounts DO NOT have access to the GPU nodes on the cluster, but you can get access by emailing Research Computing directly once a clear need arises.

Course schedule

The tentative schedule

Date	Topic
August 13	Introduction to the course & Deep Learning
August 20	Hedrick
August 27	Daniel Cortes / Deep Learning in microscopy
September 3	Khandelwal (DeepLabCut) / Student contributions
September 10	Student contributions
September 17	Student contributions
September 24	Student contributions
October 1	Student contributions
October 8	Student contributions
October 15	Student contributions
October 22	Student contributions
October 29	Student contributions
November 5	Student contributions
November 12	Last day of class – Summary discussion
November 19, 12pm	Final exam

Grading

In this course you'll be graded on your participation during class discussions (20%), your individual contribution (60%, see below) and the final exam (20%). The final exam structure is not set yet, but I don't anticipate writing a difficult exam. Furthermore, there is no curve for the course and I would be delighted to award 100% of the students A or high pass grades if your work supports that outcome.

Student contributions

Student contributions are, simply put, student-authored contributions to the general seminar topic. Each contribution will include a 15 minute presentation to the class followed by 15 minutes of discussion and be backed up by written content added to the Sakai course Wiki with a first set of revisions to that Wiki entry after review by the instructor and other students. The actual contribution content and topics may vary widely. For example, your contribution could be one of:

- identifying, reading, and summarizing a recent research paper using Deep Learning in Biology
- attempting to re-create some part of the results from a research paper using its methods and data – note that failure is an option, and your experiences in trying to re-create something are well worth presenting and discussing even if (or especially if!) they don't entirely succeed
- summarizing your attempts to employ Deep Learning in your own research
- creating a primer on a sub-topic within Deep Learning, e.g. on the different types of loss functions and when they're appropriate
- something else – send me your ideas!

Covid-19

Our current understanding of Covid-19 infection suggests that prolonged indoor contact is one of the easiest ways for the virus to spread. For this reason I ask that everyone physically present in class maximize their physical distance from one another and wear a mask covering their nose and mouth. You're also welcome to switch between virtual and physical attendance; I will do my best to make the virtual experience as complete as possible.

Guidelines for discussion and interaction

Deep learning in Biology is an interdisciplinary topic and while we have come together over a shared interest in the topic each of us has a different background and different expertise to share. I do not expect everyone to be equally knowledgeable about different areas of Biology, Mathematics or Programming but everyone's contribution is equally welcome, and I expect in class discussion to be civil, courteous, and respectful. A class of 20 students spread between physical and virtual environments is too large for us to signal one another that we wish to speak, so signal me and wait for me to call on you.

Furthermore, I am committed to creating an inclusive environment in which all students are respected and valued. I will not tolerate disrespectful language or behavior on the basis of age, ability, color/ethnicity/race, gender identity/expression, marital/parental status, military/veteran's status, national origin, political affiliation, religious/spiritual beliefs, sex, sexual orientation, socioeconomic status or other visible or non-visible differences.

Honor code

Application of the honor code can be difficult to interpret in a collaborative classroom environment, and while I expect you all to learn from one another I also expect that your contribution to the class will be based on your own effort. I expect you to be making use of many online resources; this is 100% OK, just make sure that you mention and provide links back to the original material, even if it's just a Stackexchange question and answers!