

T81-558: Applications of Deep Neural Networks

Washington University in St. Louis

Section 2. Spring 2020, Monday, 2:30 PM - 5:20 PM Online & Cupples I / 215

Section 1. Spring 2020, Monday, 6:00 PM - 9:00 PM Online & Cupples I / 215

Jeff Heaton, Ph.D., jtheaton@wustl.edu

<http://sites.wustl.edu/jeffheaton>

This syllabus is subject to change as the semester progresses.

Course Description: Deep learning is a group of exciting new technologies for neural networks. Through a combination of advanced training techniques and neural network architectural components, it is now possible to create neural networks that can handle tabular data, images, text, and audio as both input and output. Deep learning allows a neural network to learn hierarchies of information in a way that is like the function of the human brain. This course will introduce the student to classic neural network structures, Convolution Neural Networks (CNN), Long Short-Term Memory (LSTM), Gated Recurrent Neural Networks (GRU), General Adversarial Networks (GAN) and reinforcement learning. Application of these architectures to computer vision, time series, security, natural language processing (NLP), and data generation will be covered. High Performance Computing (HPC) aspects will demonstrate how deep learning can be leveraged both on graphical processing units (GPUs), as well as grids. Focus is primarily upon the application of deep learning to problems, with some introduction to mathematical foundations. Students will use the Python programming language to implement deep learning using Google TensorFlow and Keras. It is not necessary to know Python prior to this course; however, familiarity of at least one programming language is assumed. This course will be delivered in a hybrid format that includes both classroom and online instruction.

Prerequisite(s): None; however, general programming experience is assumed. The Python programming language will be used for this class and reviewed as appropriate. Elements from mathematics (generally at a Calculus I level) will be introduced and explained.

Credit Hours: 3

Text: No required text.

Course Objectives:

At the completion of this course, students will be able to:

1. Explain how neural networks (deep and otherwise) compare to other machine learning models.
2. Determine when a deep neural network would be a good choice for a particular problem.
3. Demonstrate their understanding of the material through a final project.

Grade Distribution:

Class Participation/Attendance	10%
Programming Assignments (10)	50%
Kaggle Project	20%
Final Project	20%

Letter Grade Distribution:

100	A+	77	C+
94	A	75	C
87	A-	73	C-
85	B+	<73	F
82	B		
79	B-		

Course Policies:

- **General**

- **No makeup quizzes or exams will be given.**
- Grades in the **C** range represent performance that **is below expectations**; Grades in the **B** range represent performance that **meets expectations**; Grades in the **A** range represent work that is **excellent**.
- Grades will be maintained online. Students are responsible for tracking their progress by referring to the the university's Blackboard system.

- **Programs and Assignments**

- Students are expected to work independently. **Offering** and **accepting** solutions from others is an act of **plagiarism**, which is a serious offense and **all involved parties will be penalized according to the Academic Honesty Policy**. Discussion amongst students is encouraged, but when in doubt, direct your questions to the professor, tutor, or lab assistant.
- Programs that fail to compile will not receive a grade higher than 79%. Programs that do not produce the correct output will not receive a grade higher than 89%. Original programs that produce the expected output will always receive 100%. For multi-question programming assignment, each question is evaluated individually.
- **No late assignments will be accepted under any circumstances.**

- **Attendance and Absences**

- This is a hybrid format class. Attendance of the 4 inclass sessions is very important. Attendance is expected and will be taken each class. For each unexcused absence the student will loose one percentage point of their grade (e.g. 95% to 94% for one missed class).
- Students are responsible for all missed work, regardless of the reason for absence. It is also the absentee's responsibility to get all missing notes or materials.

Academic Honesty Policy Summary:

Introduction

In addition to skills and knowledge, COLLEGE/UNIVERSITY aims to teach students appropriate Ethical and Professional Standards of Conduct. The Academic Honesty Policy exists to inform students and Faculty of their obligations in upholding the highest standards of professional and ethical integrity. All student work is subject to the Academic Honesty Policy. Professional and Academic practice provides guidance about how to properly cite, reference, and attribute the intellectual property of others. Any attempt to deceive a faculty member or to help another student to do so will be considered a violation of this standard.

Instructor's Intended Purpose

The student's work must match the instructor's intended purpose for an assignment. While the instructor will establish the intent of an assignment, each student must clarify outstanding questions of that intent for a given assignment.

Unauthorized/Excessive Assistance

The student may not give or get any unauthorized or excessive assistance in the preparation of any work.

Authorship

The student must clearly establish authorship of a work. Referenced work must be clearly documented, cited, and attributed, regardless of media or distribution. Even in the case of work licensed as public domain or Copyleft, (See: <http://creativecommons.org/>) the student must provide attribution of that work in order to uphold the standards of intent and authorship.

Declaration

Online submission of, or placing one's name on an exam, assignment, or any course document is a statement of academic honor that the student has not received or given inappropriate assistance in completing it and that the student has complied with the Academic Honesty Policy in that work.

Consequences

An instructor may impose a sanction on the student that varies depending upon the instructor's evaluation of the nature and gravity of the offense. Possible sanctions include but are not limited to, the following: (1) Require the student to redo the assignment; (2) Require the student to complete another assignment; (3) Assign a grade of zero to the assignment; (4) Assign a final grade of "F" for the course. A student may appeal these decisions according to the Academic Grievance Procedure. (See the relevant section in the Student Handbook.) Multiple violations of this policy will result in a referral to the Conduct Review Board for possible additional sanctions. The full text of the Academic Honesty Policy is in the *Student Handbook*.

Accommodations Based Upon Sexual Assault

The University is committed to offering reasonable academic accommodations to students who are victims of sexual assault. Students are eligible for accommodation regardless of whether they seek criminal or disciplinary action. Depending on the specific nature of the allegation, such measures may include but are not limited to: implementation of a no-contact order, course/classroom assignment changes, and other academic support services and accommodations. If you need to request such accommodations, please direct your request to Kim Webb (kim_webb@wustl.edu), Director

of the Relationship and Sexual Violence Prevention Center. Ms. Webb is a confidential resource; however, requests for accommodations will be shared with the appropriate University administration and faculty. The University will maintain as confidential any accommodations or protective measures provided to an individual student so long as it does not impair the ability to provide such measures.

If a student comes to me to discuss or disclose an instance of sexual assault, sex discrimination, sexual harassment, dating violence, domestic violence or stalking, or if I otherwise observe or become aware of such an allegation, I will keep the information as private as I can, but as a faculty member of Washington University, I am required to immediately report it to my Department Chair or Dean or directly to Ms. Jessica Kennedy, the University's Title IX Director. If you would like to speak with directly Ms. Kennedy directly, she can be reached at (314) 935-3118, jwkennedy@wustl.edu, or by visiting the Title IX office in Umrath Hall. Additionally, you can report incidents or complaints to the Office of Student Conduct and Community Standards or by contacting WUPD at (314) 935-5555 or your local law enforcement agency. See: Title IX

You can also speak confidentially and learn more about available resources at the Relationship and Sexual Violence Prevention Center by calling (314) 935-8761 or visiting the 4th floor of Seigle Hall. See: RSVP Center

Bias Reporting

The University has a process through which students, faculty, staff and community members who have experienced or witnessed incidents of bias, prejudice or discrimination against a student can report their experiences to the University's Bias Report and Support System (BRSS) team. See: <http://brss.wustl.edu>

Mental Health

Mental Health Services professional staff members work with students to resolve personal and interpersonal difficulties, many of which can affect the academic experience. These include conflicts with or worry about friends or family, concerns about eating or drinking patterns, and feelings of anxiety and depression. See: shs.wustl.edu/MentalHealth

Course Outline:

The weekly coverage might change as it depends on the progress of the class (note: for this class, weeks end Sunday evening the day before class). All assignments are due by midnight on the date specified, except the Kaggle submitted project (which is in GMT, see below).

Module	Date	Content
Module 1	Meet on: 01/13/2020	<ul style="list-style-type: none">• Python Preliminaries• We will meet on campus this week! (1st Meeting)
Module 2	Week of 01/27/2020	<ul style="list-style-type: none">• Python for Machine Learning• Module 1 Assignment due 01/28/2020
Module 3	Week of 02/03/2020	<ul style="list-style-type: none">• TensorFlow and Keras for Neural Networks• Module 2 Assignment due 02/04/2020
Module 4	Week of 02/10/2020	<ul style="list-style-type: none">• Training for Tabular Data• Module 3 Assignment due 02/11/2020
Module 5	Week of 02/17/2020	<ul style="list-style-type: none">• Regularization and Dropout• Module 4 Assignment due 02/18/2020
Module 6	Meet on: 02/24/2020	<ul style="list-style-type: none">• Convolutional Neural Networks (CNNs) for Vision• Module 5 Assignment due 02/25/2020• We will meet on campus this week! (2nd Meeting)
Module 7	Week of: 03/02/2020	<ul style="list-style-type: none">• Generative Adversarial Networks (GANs)• Module 6 Assignment due 03/03/2020
Module 8	Week of 03/16/2020	<ul style="list-style-type: none">• Kaggle and Advanced Data Sets• Module 7 Assignment due 03/17/2020
Module 9	Meet on: 03/23/2020	<ul style="list-style-type: none">• Transfer Learning• Module 8 Assignment due 03/24/2020• We will meet on campus this week! (3rd Meeting)
Module 10	Week of 03/30/2020	<ul style="list-style-type: none">• Time Series in Keras• Module 9 Assignment due 03/31/2020
Module 11	Week of 04/06/2020	<ul style="list-style-type: none">• Natural Language Processing• Module 10 Assignment due 04/07/2020
Module 12	Week of 04/13/2020	<ul style="list-style-type: none">• Reinforcement Learning• Kaggle Assignment due: 04/13/2020 (approx 4-6PM, due to Kaggle GMT timezone)
Module 13	Meet on: 04/20/2020	<ul style="list-style-type: none">• Deployment and Monitoring• Top Kaggle presentations by students• We will meet on campus this week! (4th Meeting)
Module 14	Week of 04/27/2020	<ul style="list-style-type: none">• Other Neural Network Techniques• Student presentations on Zoom• Final Project due 05/04/2020