Course ID and Title: CSCI 699: Advanced Topics in Deep Learning
Units: 4
Term—Day—Time: Fall 2020

IMPORTANT:
The general formula for contact hours is as follows:
Courses must meet for a minimum of one 50-minute session per unit per week over a 15-week semester. Standard fall and spring sessions (001) require a final summative experience during the University scheduled final exam day and time.

(Please refer to the *Contact Hours Reference* guide.)

Location: Online

Instructor:
Office:
https://usc.zoom.us/j/98169788926?pwd=M1QwOUUyNnowODlLWGJGbzlSQkdzQT09
Office Hours: Friday 11am-noon
Contact Info: Email: yanliu.cs@usc.edu (the instructor will reply within 24 hours)

Teaching Assistant:
Office: Physical or virtual address
Office Hours:
Contact Info: Email, phone number (office, cell), Skype, etc.

IT Help: Group to contact for technological services, if applicable.
Hours of Service:
Contact Info: Email, phone number (office, cell), Skype, etc.

Course Description

Deep learning and artificial intelligence is deemed as one of the most important revolutions in computer science in the past decade. It has been demonstrated to be extremely effective in learning and prediction tasks, such as computer vision, natural language processing, robotics and so on. The development in the field is so fast (with hundreds of papers in one week) that it becomes rather hard for individual researchers and students to catch up. Therefore a class covering recent advances in deep learning is extremely timely and important.

The course aims to introduce recent important advances in deep learning models, such as deep reinforcement learning, meta-learning, Generative Adversarial Networks (GAN), Variational Autoencoders, graph neural networks and interpretation of neural networks. Enrolled students should have basic knowledge in machine learning and deep learning, basic skills in conduct independent research, and programming skills. In the class, the students will have guided discussions on recent papers of advance topics in deep and conduct course project to utilize the knowledge discussed in the class.

At the end of the course, the students are expected to be able to do the following: (1) understanding the mathematical formulation of different types of deep learning models; (2) apply deep learning models to real-world applications; (3) developing novel deep learning models for applications of their interest, preferably publishable in top conferences in machine learning.

Learning Objectives and Outcomes

Students will learn fundamental knowledge and gain hands-on experience in order to:

1. Applying existing machine learning and deep learning models
2. Design new machine learning models to address novel challenges in practical applications

Prerequisite(s): CSCI 567; basic knowledge in statistics and probability, linear algebra, optimization. Basic programming skills are required.**Co-Requisite(s):** course(s) that must be taken prior to or simultaneously

Concurrent Enrollment: N/A

Recommended Preparation: N/A

Course Notes

Grading Type: Letter grade.

Technological Proficiency and Hardware/Software Required

Basic programming skills are required.

Required Readings and Supplementary Materials

- Ian Goodfellow and Yoshua Bengio and Aaron Courville (2016). Deep Learning, MIT Press.
- Reading list is provided below in details

Description and Assessment of Assignments

The learning outcome will be assessed through in-class presentations, reading reviews and class project.

1) Project requirements

The purpose of the project is to practice the concepts learned in the class and apply them to address an important application of the students' interest. The project will be conducted by a group of two students.

Timeline:

Aug 24 – Sep 4: Identifying team members and project topics

Sep 11: Proposal due (summarizing team member, topics and milestone)

Oct 23: Mid-term report due (discussing data description, preliminary results)

Nov 20: Project presentation and Poster session (open to all faculty and students)

Dec 7 : Final report due (describing the task, proposed solutions, major discovery, lessons learned)

Sample projects “Deep Learning for analyzing misinformation on twitter data”:
In this project, students will develop effective topic models for twitter data. Students can easily find resources available online, including twitter API and the C++ topic modeling code (e.g. LDA model). A typical team size of this project is 2 (persons). The team will work together on collecting misinformation on Twitter, examining the preliminary results, identifying one challenge in current deep learning solutions, and providing a reasonable solution.

Grading breakdown of the course project:
Proposal: 5%
Mid-term report: 5%
Final report: 15%
Presentation: 20%
Poster: 5%

All members in one team will get the same grade

2) Project Presentation
The expected presentation of each paper is 30 minutes. The students should develop their own slides, but they are allowed to use materials online for part of their slides (no more than 50%).

Grading Breakdown

Including the above detailed assignments, how will students be graded overall? Participation should be no more than 15%, unless justified for a higher amount. All must total 100%.

Assignment	**Points**	**% of Grade**
Reading review		20
In-class presentation		30
Class project		50
TOTAL		100

Grading Scale (Example)

Course final grades will be determined using the following scale

A	95-100
A-	90-94
B+	87-89
B	83-86
B-	80-82
C+	77-79
C	73-76
C-	70-72
D+	67-69
D	63-66
D-	60-62
F	59 and below

Assignment Rubrics
Include assignment rubrics to be used, if any.

Assignment Submission Policy
Each of the review assignments is to be submitted in two weeks, as specified in the course schedule.

Grading Timeline
Each assignment is graded and returned with feedback in one week.

Additional Policies
N.A.

Course Schedule: A Weekly Breakdown

(Please refer to the *Contact Hours Reference guide*.)

Week 1 Course overview and fundamental maths for ML (Aug 28th)	
Course overview, syllabus, projects, grading	
MDPs, Value Iteration, Policy Iteration, Function approximation, Q-learning	
Graphs, Adjacency matrix, Laplacian, Eigenspectrum of laplacian, Convolution on graphs	
Probability densities, divergence metrics, density transformation via jacobian	
Week 2 Deep Reinforcement Learning (Sep 4th)	
Human-level control through deep reinforcement learning	https://www.nature.com/articles/nature14236
Asynchronous Methods for Deep Reinforcement Learning	https://arxiv.org/abs/1602.01783
Imagination-Augmented Agents for Deep Reinforcement Learning	https://arxiv.org/abs/1707.06203
Graph Networks as Learnable Physics Engines for Inference and Control	https://arxiv.org/pdf/1806.01242
Week 3 Meta-learning (Sep 11th)	
Learning to Optimize	https://arxiv.org/abs/1606.01885
Model-Agnostic Meta-Learning for Fast Adaptation of Deep Networks	https://arxiv.org/abs/1703.03400
Continuous Adaptation via Meta-Learning in Nonstationary and Competitive Environments	https://arxiv.org/abs/1710.03641
Learning to Generalize: Meta-Learning for Domain Generalization	https://arxiv.org/pdf/1710.03463.pdf
Week 4 Project idea discussion and proposal (Sep 18th)	
Week 5 GANs (Sep 25th)	
NIPS 2016 Tutorial: Generative Adversarial Networks	https://arxiv.org/abs/1701.00160
Tempered Adversarial Networks	http://proceedings.mlr.press/v80/sajjadi18a.html
RadialGAN: Leveraging multiple datasets to improve target-specific predictive models using Generative Adversarial Networks	http://proceedings.mlr.press/v80/yoon18b/yoon18b.pdf
The relativistic discriminator: a key element missing from standard GAN	https://openreview.net/forum?id=S1erHoR5t7

Week 6 VAEs and Normalizing flows – part I (Oct 2nd)	
Tutorial on Variational Autoencoders	https://arxiv.org/abs/1606.05908
Variational inference with normalizing flows	https://arxiv.org/abs/1505.05770
Week 7 VAEs and Normalizing flows – part II (Oct 9th)	
Improving Variational Inference with Inverse Autoregressive Flow	https://papers.nips.cc/paper/6581-improved-variational-inference-with-inverse-autoregressive-flow
Density estimation using real NVP	https://arxiv.org/abs/1605.08803
Glow: Generative Flow with Invertible 1x1 Convolutions	http://papers.nips.cc/paper/8224-glow-generative-flow-with-invertible-1x1-convolutions
Week 8: Learning on graphs - part I (Oct 16th)	
Graph Signal Processing: Overview, Challenges and Applications	https://arxiv.org/abs/1712.00468
Convolutional Networks on Graphs for Learning Molecular Fingerprints	https://arxiv.org/abs/1509.09292
Gated Graph Sequence Neural Networks	https://arxiv.org/abs/1511.05493
Semi-Supervised Classification with Graph Convolutional Networks	https://arxiv.org/abs/1609.02907
Week 9 Learning on graphs - part II (Oct 23th)	
Inductive Representation Learning on Large Graphs	https://arxiv.org/abs/1706.02216
Hierarchical Graph Representation Learning with Differentiable Pooling	https://arxiv.org/pdf/1806.08804
GraphRNN: Generating Realistic Graphs with Deep Auto-regressive Models	https://arxiv.org/pdf/1802.08773
Learning Graphical State Transitions	https://openreview.net/forum?id=HJ0NvFzxl
Week 10 Relational networks - part I (Oct 30th)	
Interaction Networks for Learning about Objects, Relations and Physics	https://papers.nips.cc/paper/6418-interaction-networks-for-learning-about-objects-relations-and-physics
A simple neural network module for relational reasoning	https://papers.nips.cc/paper/7082-a-simple-neural-network-module-for-relational-reasoning
Neural Message Passing for Quantum Chemistry	https://arxiv.org/abs/1704.01212
Graph Attention Networks	https://arxiv.org/abs/1710.10903
week 11 Relational networks - part II (Nov 6th)	

Relational inductive biases, deep learning, and graph networks	https://arxiv.org/pdf/1806.01261
Relational Neural Expectation Maximization: Unsupervised Discovery of Objects and their Interaction	https://arxiv.org/abs/1802.10353
Neural Relational Inference for Interacting Systems	https://arxiv.org/abs/1802.04687
End-to-End Differentiable Physics for Learning and Control	https://papers.nips.cc/paper/7948-end-to-end-differentiable-physics-for-learning-and-control
Week 12 Interpretability – Part I (Nov 13th)	
"Why Should I Trust You?" Explaining the Predictions of Any Classifier	https://www.kdd.org/kdd2016/papers/files/rfp0573-ribeiroA.pdf
Interpretability Beyond Feature Attribution: Quantitative Testing with Concept Activation Vectors (TCAV)	https://arxiv.org/pdf/1711.11279.pdf
To Trust Or Not To Trust A Classifier	https://papers.nips.cc/paper/7798-to-trust-or-not-to-trust-a-classifier.pdf
Improving Simple Models with Confidence Profiles	http://papers.nips.cc/paper/8231-improving-simple-models-with-confidence-profiles.pdf
Week 13 Interpretability – Part II (Nov 20th)	
Axiomatic Attribution for Deep Networks	https://arxiv.org/pdf/1703.01365
Intelligible Models for HealthCare: Predicting Pneumonia Risk and Hospital 30-day Readmission	http://people.dbmi.columbia.edu/noemie/papers/15kdd.pdf
GradCAM: Visual Explanations from Deep Networks via Gradient-based Localization	https://arxiv.org/pdf/1610.02391.pdf
Exact and Consistent Interpretation for Piecewise Linear Neural Networks: A Closed Form Solution	https://dl.acm.org/citation.cfm?doid=3219819.3220063
Final **Project final presentation (TBD)**	

Statement on Academic Conduct and Support Systems

Academic Conduct:

Plagiarism – presenting someone else's ideas as your own, either verbatim or recast in your own words – is a serious academic offense with serious consequences. Please familiarize yourself with the discussion of plagiarism in SCampus in Part B, Section 11, "Behavior Violating University Standards" policy.usc.edu/scampus-part-b. Other forms of academic dishonesty are equally unacceptable. See additional information in SCampus and university policies on scientific misconduct, policy.usc.edu/scientific-misconduct.

Support Systems:

Student Health Counseling Services - (213) 740-7711 – 24/7 on call
engemannshc.usc.edu/counseling
Free and confidential mental health treatment for students, including short-term psychotherapy, group counseling, stress fitness workshops, and crisis intervention.

National Suicide Prevention Lifeline - 1 (800) 273-8255 – 24/7 on call
suicidepreventionlifeline.org
Free and confidential emotional support to people in suicidal crisis or emotional distress 24 hours a day, 7 days a week.

Relationship and Sexual Violence Prevention Services (RSVP) - (213) 740-4900 – 24/7 on call
engemannshc.usc.edu/rsvp
Free and confidential therapy services, workshops, and training for situations related to gender-based harm.

Office of Equity and Diversity (OED) | Title IX - (213) 740-5086
equity.usc.edu, titleix.usc.edu
Information about how to get help or help a survivor of harassment or discrimination, rights of protected classes, reporting options, and additional resources for students, faculty, staff, visitors, and applicants. The university prohibits discrimination or harassment based on the following protected characteristics: race, color, national origin, ancestry, religion, sex, gender, gender identity, gender expression, sexual orientation, age, physical disability, medical condition, mental disability, marital status, pregnancy, veteran status, genetic information, and any other characteristic which may be specified in applicable laws and governmental regulations.

Bias Assessment Response and Support - (213) 740-2421
studentaffairs.usc.edu/bias-assessment-response-support
Avenue to report incidents of bias, hate crimes, and microaggressions for appropriate investigation and response.

The Office of Disability Services and Programs - (213) 740-0776
dsp.usc.edu
Support and accommodations for students with disabilities. Services include assistance in providing readers/notetakers/interpreters, special accommodations for test taking needs, assistance with architectural barriers, assistive technology, and support for individual needs.

USC Support and Advocacy - (213) 821-4710
studentaffairs.usc.edu/ssa
Assists students and families in resolving complex personal, financial, and academic issues adversely affecting their success as a student.
Diversity at USC - (213) 740-2101
diversity.usc.edu
Information on events, programs and training, the Provost's Diversity and Inclusion Council, Diversity Liaisons for each academic school, chronology, participation, and various resources for students.

USC Emergency - UPC: (213) 740-4321, HSC: (323) 442-1000 – 24/7 on call
dps.usc.edu, emergency.usc.edu
Emergency assistance and avenue to report a crime. Latest updates regarding safety, including ways in which instruction will be continued if an officially declared emergency makes travel to campus infeasible.

USC Department of Public Safety - UPC: (213) 740-6000, HSC: (323) 442-120 – 24/7 on call
dps.usc.edu
Non-emergency assistance or information.