CSCI 699: Dynamics of Representation Learning

Units: 4

Spring 2022—Mon/Wed—12:00-1:50pm

Location: GFS 223

Course website:
https://sites.google.com/view/699dynamicsofrep/

Instructor: Greg Ver Steeg

Office: https://usc.zoom.us/j/6818370261

Office Hours: by appointment

Contact Info: gregv@isi.edu, cell: 626-840-5901.

Teaching Assistant: Umang Gupta

Office: Physical or virtual address

Office Hours:

Contact Info: umanggup@usc.edu

No person steps in the same river twice, for it's not the same river and not the same person.
-Heraclitus (500 BC, after being asked whether the test set will resemble the training set)

Course Description

Deep learning is a swiftly changing field, and many of its key ingredients can themselves be viewed as dynamical processes. Optimizers used in deep learning, for instance, exploit dynamics that lead to solutions with desirable properties that are not intrinsic to the loss function itself. Dynamics appears in many other guises. Sampling relies on MCMC dynamics and finds applications in Bayesian optimization, latent factor inference, and energy-based probability modeling. Learning tasks must be dynamic to reflect a world that is constantly changing. Deep architectures and normalizing flows can be viewed as dynamical processes that transform data over time through a sequence of layers. This course will be research-oriented and survey recent papers in the field. Rather than focusing on state-of-the-art numbers on common benchmark tasks, we will look for common mathematical threads and try to develop deeper intuition about representation learning through the lens of dynamics.

Learning Objectives

The goal of this course is to develop a deeper appreciation of common deep learning ideas such as SGD, density models, or transfer learning by looking at them in terms of dynamics, equilibria, and Bayesian reasoning. Deep learning is rife with tricks that practitioners can try in order to build good models. However, a PhD researcher should not be content with blind trial and error or speculation about how and why something works, but instead should be constantly striving to understand deeper principles. The hope is that by the end of this course, students will be familiar with some of the theoretical frameworks which are being applied to understand deep learning successes and failures, and equipped to understand and contribute new insights in the field.

Prerequisite(s):

(1) General proficiency in computer science (linear algebra, calculus, probability)
(2) Basic programming including some familiarity with either Tensorflow or PyTorch.

Course Notes

TBD: To accommodate the possibility of hybrid attendance (in-person and remote), this course will be taught on campus in a classroom that also supports synchronous online attendance.
Copies of the lecture slides and other class information will be posted on the course website.
Students will be graded on course presentations and projects.

Required Readings and Supplementary Materials

Students are expected to study recent research papers, with starting points listed on the syllabus. Background and context for certain topics can be found in two open source deep learning books: Deep Learning Book and the more recent Dive Into Deep Learning, referenced in the Syllabus as DLB Ch. X or DIDL Ch. X.

Description and Assessment of Assignments

Each project will include reproduction of some common baseline task followed by an exploration of some distinct variation on that task. Whether projects are individual or teams is to be determined, and may take into account student feedback. The themes for each project follow, and more specific instructions for each will be given in class.
Project 1: Optimization dynamics and loss landscapes
Project 2: Sampling dynamics
Project 3: Dynamic tasks

Write-ups

Write-ups should be submitted as LaTeX generated PDFs using the NeurIPS style file. They should at least contain an abstract, introduction section giving motivation and a few citations to relevant work, and a results section. For each project, I'll ask for specific types of results to be included. Grading will be based on readability, inclusion of all relevant details to understand the experiments, accurate reproduction of the baseline task, and including specified result types for the task variation explored. Grades *will not* depend on getting a positive or negative result. I will encourage students to consider ideas that are risky and therefore likely to fail. There's no page limit, but I would guess that around 4 pages in single column format are required. Don't be stingy with plot sizes, and be sure to include legible axis labels and legends.

Code
I'm fluent in PyTorch and conversant in Tensorflow. The main thing I'm looking for here is that the code is well-organized and commented, and matches what was used to produce the results. I like IPython notebooks, or traditional commented python code.

Presentations
In class presentations will present the work and include time for questions and discussion. The exact length allotted will depend on class and team size.

Grading Breakdown

Including the above detailed assignments, how will students be graded overall? Participation should not exceed 15% of the total grade. Where it does, the syllabus must provide an added explanation. No portion of the grade may be awarded for class attendance but non-attendance can be the basis for lowering the grade, when clearly stated on the syllabus. The sum of percentages must total 100%.

Assessment Tool (assignments)	**Points**	**% of Grade**
Project 1 write-up	10	2/10
Project 1 presentation	5	1/10
Project 2 write-up	10	2/10
Project 2 presentation	5	1/10
Project 3 write-up	10	2/10
Project 3 presentation	5	1/10
Acceptable code included for all projects	5	1/10
TOTAL	50	100%

Grading Scale

Course final grades will be determined using the following scale

A	90-100
A-	80-90
B+	75-80
B	70-75
B-	65-70
C+	60-65
C	50-60
F	Below 55

Assignment Submission Policy

By email, before 11:59pm of the due date.

Grading Timeline

Assignments will be graded within one week after the due date.

Additional Policies

Late homework policy: you are given 4 late days for the assignments and project proposal/survey (no late days for the final project report), to be used in integer amounts and distributed as you see fit. Additional late days will each result in a deduction of 10% of the grade of the corresponding assignment.

Course Schedule: A Weekly Breakdown

Suggested papers are subject to change – the latest version will be on the course website: https://sites.google.com/view/699dynamicsofrep/

	Topics/Daily Activities	**Readings/Preparation**	**Deliverables**
Week 1 Jan 10	Class Intro and math review: connecting continuous & discrete dynamics, differential equations, classical dynamics, Markov chains Optimization dynamics	An overview of gradient descent optimization algorithms DIDL Ch. 11	
Week 2 Jan 17	**MLK Day** (Monday) Visualizing loss landscapes, dynamics, and solutions	Visualizing the loss landscape Understanding generalization through visualization https://losslandscape.com/faq/	
Week 3 Jan 24	Bayesian statistics review Epistemic and aleatoric uncertainty Bayesian model averaging	Deep Ensembles SWA, SWAG, Multi SWAG Dataset cartography	
Week 4 Jan 31	Mon: discuss projects Wed: Theory of over-parametrized deep learning	Neural tangent kernel Opt. landscape of over-parametrized shallow networks	
Week 5 Feb 7	Research and communication dynamics: writing as story-telling, staying skeptical	Troubling trends in ML Mythos of model interpretability Abstract writing notes	**Project 1 write-ups and code** Due 11:59 pm Feb. 10
Week 6 Feb 14	Project 1 Presentations		**Project 1 presentations**
Week 7 Feb 21	**President's day** (Mon) Wed: Sampling applications in machine learning overview: Energy models, latent factor models, Bayesian inference	A complete recipe for stochastic gradient MCMC Stoch. Grad. Langevin Dynamics	
Week 8 Feb 28	Sampling: Langevin dynamics, MCMC, HMC Importance sampling	Radford Neal long classics: HMC & AIS	

	Annealed importance sampling		
Week 9 March 7	Mon: Training energy models via sampling Wed: Discuss projects	How to train your energy based model, JEM	
March 14	**Spring break – no class**		
Week 10 March 21	Task dynamics: Industry problems and relevance, Transfer learning Domain adaptation, sub-population shift, Causal and anti-causal shift, Meta-learning	How to train your MAML Simple fine-tuning FTW (2) Causal and anti-causal shift WILDS domain shift tasks Can you trust your model's uncertainty?	**Project 2 write-up and code** Due 11:59 pm March 23
Week 11 March 28	**Project 2 presentations**		**Project 2 presentations**
Week 12 April 4	Domain generalization Out-of-distribution detection vs generalization	Lost domain generalization Invariant risk minimization OOD hierarchy of features Density models not sufficient for OOD Anomaly review	
Week 13 April 11	Defining functions in terms of fixed point dynamics, optimizing with implicit differentiation Flows as dynamics, neural ODE	Invertible ResNets Deep equilibrium models Implicit deep learning Out-of-the-box idea: nonlinear parallel equation solving instead of feedforward Neural ODE	
Week 14 April 18	Equilibrium, Stochastic normalizing flows	Original nonequilibrium motivation Denoising diffusion (2) Stochastic normalizing flows	
Week 15 April 25	Students' choice on related topics: reinforcement learning, minimax optimization/GANs, Differentiable physics, message-passing in graph neural networks, continual learning…		**Project 3 write-up and code** Due 11:59 pm April 27
FINAL	**Project 3 presentations**		Refer to the final exam schedule in the USC *Schedule of Classes* at classes.usc.edu.

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." Other forms of academic dishonesty are equally unacceptable. See additional information in SCampus and university policies on scientific misconduct.

Support Systems:

Counseling and Mental Health
phone number (213) 740-9355
On call 24/7
Free and confidential mental health treatment for students, including short-term psychotherapy, group counseling, stress fitness workshops, and crisis intervention.

National Suicide Prevention Lifeline
Phone number 1 (800) 273-8255
On call 24/7
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)
Phone Number (213) 740-9355(WELL), press "0" after hours
On call 24/7
Free and confidential therapy services, workshops, and training for situations related to gender-based harm.

USC Office of Equity, Equal Opportunity, and Title IX
Phone number (213) 740-5086
Title IX Office (213) 821-8298
Information about how to get help or help someone affected by harassment or discrimination, rights of protected classes, reporting options, and additional resources for students, faculty, staff, visitors, and applicants.

Reporting Incidents of Bias or Harassment
Phone number (213) 740-5086 or (213) 821-8298
Avenue to report incidents of bias, hate crimes, and microaggressions to the Office of Equity, Equal Opportunity, and Title IX for appropriate investigation, supportive measures, and response.

The Office of Disability Services and Programs
Phone number (213) 740-0776
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 Campus Support and Intervention
Phone number (213) 821-4710
Assists students and families in resolving complex personal, financial, and academic issues adversely affecting their success as a student.

[Diversity at USC](#)
Phone number (213) 740-2101
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 phone number (213) 740-4321
HSC phone number (323) 442-1000
On call 24/7
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 phone number (213) 740-6000
HSC phone number (323) 442-1200
On call 24/7
Non-emergency assistance or information.