DSCI-565:
Introduction to Deep Learning for Data Science

Units: 4

Instructor: Ke-Thia Yao
Contact Info: kyao@isi.edu

Course Description
Deep learning is a branch of machine learning focusing on artificial neural networks with multiple layers of neurons, and deep learning technology is behind many recent advances in artificial intelligence, including voice recognition, image comprehension, natural language understanding and self-driving cars. This course will provide students with a solid understanding of core deep learning algorithms and practical experience in building and applying deep learning networks through homework assignments, programming assignments using Jupyter notebooks, midterm and a final class project.

Following the philosophy of the course textbook Dive into Deep Learning this course adopts an learning-by-doing method, which involves first introducing deep learning concepts and then providing exercises on real datasets. DSCI 565 takes a practitioner's approach with emphasis on intuition about how to use deep learning algorithms with different types of data for different domains.

Learning Objectives
After successful completion of this course, students will be able to:

- Understand the principles of deep learning and deep neural networks
- Know the suitability of specific deep learning algorithms to various data domains, such as images, text, and graphs
- Design, implement and train deep neural models
- Perform regularization, training optimization and hyperparameter on deep learning models

Requirements
Pre-requisites for this course are DSCI 510 (Priciples of Programming for Data Science) and DSCI 552 (Machine Learning for Data Science).

Readings
The textbook for the course is:

Zhang, Aston and Lipton, Zachary C. and Li, Mu and Smola, Alexander J. (2021). "Dive into Deep Learning." arXiv preprint arXiv:2106.11342

It is freely available online along with materials at https://d2l.ai/. This textbook provides up to date deep learning topics, including attention mechanisms, transformers and generative adversarial networks. All the codes in this book are runnable under Jupyter Notebook Multiple implementations of the algorithms are provided, based on three popular deep learning frameworks including PyTorch, MXNet, and TensorFlow.

Supplementary reading material for this course will include academic papers and other reference books, including Deep Learning (Goodfellow, Bengio and Courville, 2016), Deep Learning with PyTorch (Antiga, Stevens and Viehmann, 2020) and Graph Representation Learning (Hamilton, 2020).

In addition, reading assignments will be selected from various periodicals and other sources.

Grading Breakdown

Assessment Tool (assignments)	% of Grade
Midterm Exam (Week 9)	30%
Homework/Programming Assignments	35%
Class Participation	5%
Semester Project	30%
TOTAL	100%

Assignment Submission Policy
Assignments and semester project will be submitted electronically via Blackboard. Assignments will be accepted after the deadline with the following grade penalties. Cumulative of 10% times number of days late:

- 1 day late: lose 10%
- 2 days late: lose 30% (10% + 20%)
- 3 days late: lose 60% (30% + 30%)
- Greater than 4 days late not accepted

No personal emergencies will be entertained (with the exception of the USC granted emergencies, in which case official documents need to be shown).

Participation
Students are expected to actively participate in this course. Participation includes:

- Careful reading and viewing of assigned materials by the date due
- Regular, substantive contributions to discussions and in-class questions
- Active engagement with online content

Course grades for students who do not contribute to the course through active participation in class may be affected.

Homework assignments
The homework assignments will be primarily based on the questions at the end of each chapter sections of the Zhang et al books, as well as additional questions relevant to course lecture content.

Programming assignments
The programming assignments will be in Jupyter Notebook format. To run the notebook students can use either Google Colab or their own computers.

- Design and implement convolutional neural networks for object classification from images
- Design and implement recurrent neural networks with attention mechanism for natural language classification
- Design and implement graph embedding models for link predictions

Semester project: the purpose of the class project is for you to practice applying deep learning algorithms and models to challenging learning problems. Machine learning conferences sometimes have dataset contribution tracks, for example NeurIPS 2022 Datasets and Benchmarks (https://nips.cc/virtual/2022/events/datasets-benchmarks-2022) and CPVR 2022 Dataset Contributions (https://cvpr2022.thecvf.com/dataset-contributions) . One method to developing a project is to find a dataset/problem of interest to you; to replicate the results of the paper; and to improve its results or to develop alternative approaches and/or alternative problems. Also, you are encouraged to identify new deep learning applications and novel ways of combining deep learning models. Working as a group is permitted if the project is large enough to justify this.

Project Timeline:

- Week 3: Identifying team members and project topics
- Week 6: Proposal due (team member, topics and milestone)
- Week 8: Mid-term report due (data description, preliminary results)
- Week 15:Project presentation and Poster session (open to all faculty and students)
- Finals week: Final report due (task and model description, major discovery, lessons learned)

Sample projects:

- Urban Forest Monitoring (https://google.github.io/auto-arborist/): the goal of this project is to develop tree genus-level classifiers for urban tree images. Students should investigate approaches to dealing with long-tail classification problems, and implement and test approaches against the dataset.
- Word and image association (https://nips.cc/virtual/2022/poster/55689, https://winogavil.github.io/): the goal of this project is to develop and to integrate vision and language deep learning models to solve word-image association problems. Students should investigate techniques, such as eliciting multiple word senses from language models and incorporating background knowledge to improve associations.

Project Presentation: The presentation should be approximately 10 minutes in duration and should cover all aspects of the project development, solution and outcomes. You may use PowerPoint or LaTeX or similar. Typically plan on one slide per minute; content should be phrases in bullets, not prose. Do not read word-for-word from your slides!

Poster: Create a poster describing your project; the poster should be suitable for presentation at a technical conference. Please use A0-sized paper (841 mm × 1189 mm = 33.1" × 46.8"). Start with a title above the name(s) and affiliation(s) of the author(s); include sections for introduction/motivation, background, methodology, results, conclusions and references; most of the section content will not be prose but bullet points; you likely will not be able to include all of the technical details, but should include graphs, photos and diagrams that make it easy for the audience to grasp the main points and impact of your work. Remember that your poster should be clearly read from a distance away (*e.g.*, the title should be roughly 25 mm = 1" high, or about 100pt font; the smallest text should typically be at least 36pt).

Final Report: The project should be documented with a written report that documents the motivation for the project, assumptions used, methodology, results, observations, conclusions and references. The main text of the report will typically be 8–10 pages (single-spaced, 1" margins, 11pt Times or 10pt Arial or similar) plus appendices (graphs, code, etc.).

Grading breakdown: All aspects of the project combined are 30% of the semester grade, with a breakdown:

- Proposal: 5%
- Mid-term report: 5%
- Final report: 5%
- Presentation: 10%
- Poster: 5%

Course Schedule

The lectures of this course are divided into two parts. The first part, approximately two-thirds of the course, focuses on the state of modern deep learning and its applications to practical problems. The second part focuses on more advanced topics in deep learning covering recent research as this is an area of very fast advances.

Week 1: Overview of deep learning and review of machine learning basics

Topics:

- Success applications of deep learning
- The essence of deep learning (many layers, end-to-end training, no need for feature engineering)
- Key components of deep learning (data, models, objective functions, optimization algorithms)
- Kinds of machine learning problems (supervised, unsupervised, self-supervised, reinforcement learning)

Readings:

- Zhang et al, Chapter 1
- (Optional) Goodfellow et al, Chapter 1
- (Optional) Goodfellow et al, Chapter 7.6, 7.7
- (Optional) Antigua et al, Chapter 1

Week 2: Deep learning preliminaries

Topics:

- Data manipulation and preprocessing (indexing, slicing, broadcasting)
- Linear algebra (matrices and tensors, tensor arithmetic, reductions, products, norms)
- Calculus (derivatives, gradients, chain rule)
- Automatic differentiation
- Probabilities and statistics (random variables, chain rule, expectation, variance and covariance, Bayes' Rule)

Readings:

- Zhang et al, Chapter 2
- (Optional) Goodfellow et al, Chapters 2 and 3
- (Optional) Antigua et al, Chapter 3

Week 3: Linear neural networks for regression and classification

Project team member identification and project topics due

Topics:
- Linear regression
- Object-oriented design (data, model, training) and implementation using PyTorch
- Softmax regression for classification and implementation using PyTorch
- Image classification

Readings:
- Zhang et al, Chapter 3
- Zhang et al, Chapter 4
- (Optional) Antiga et al, Chapter 5.1, 5.2, 5.3

Week 4: Deep feedforward networks

Topics:
- Learning XOR problem
- Multilayer perceptrons and hidden layers
- Forward and backward propagation
- Numerical Stability (vanishing and exploding gradients)
- Overfitting and regularization
- Dropout

Readings:
- Zhang et al, Chapter 5
- (Optional) Goodfellow et al, Chapter 6
- (Optional) Goodfellow et al, Chapter 7.12
- (Optional) Antiga et al, Chapter 5.4, 5.5
- (Optional) Antiga et al, Chapter 6

Week 5: Deep learning builder's guide

Topics:
- Layers and modules
- Parameter management
- Parameter initialization and lazy initialization
- Custom layers
- GPUs

Readings:
- Zhang et al, Chapter 6

Week 6: Convolutional neural networks

Project Proposal Due

Topics:
- Convolutional operator and its motivation
- Padding and stride
- Pooling
- Multiple input and multiple output channels
- Image classification and LeNet as an example
- Batch normalization

- Designing convolution networks
- Modern convolutional neural networks: AlexNet, VGG, NiN, GoogLeNet, ResNet, DenseNet

Readings:

- Zhang et al, Chapter 7
- Zhang et al, Chapter 8
- (Optional) Zhang et al, Chapter 14
- (Optional) Goodfellow et al, Chapter 9
- (Optional) Antiga et el, Chapter 4.1, 4.2
- (Optional) Antiga et al, Chapter 8

Week 7: Recurrent neural networks

Topics:

- Sequence and text problems and models
- Unfolding computational graphs
- Recurrent neural networks with and without hidden states
- Backpropagation through time
- Long-term dependency challenge
- Modern recurrent neural networks: LSTM, GRU, deep recurrent neural networks, bidirectional recurrent neural networks, encoder-decoder neural networks, and seq2seq machine translation.

Readings:

- Zhang et al, Chapter 9
- Zhang et al, Chapter 10
- (Optional) Goodfellow et al, Chapter 10
- (Optional) Antiga et al, Chapter 4.4, 4.5

Week 8: Attention mechanisms and transformers

Project Mid-term report due

Topics:

- Attention cues
- Attention pooling
- Multi-head attention, self-attention and positional encoding.
- The transformer architecture.

Readings:

- Zhang et al, Chapter 11
- Vaswani, Ashish, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. “Attention Is All You Need.” In *Advances in Neural Information Processing Systems*. Vol. 30. Curran Associates, Inc. https://proceedings.neurips.cc/paper/2017/hash/3f5ee243547dee91fbd053c1c4a845aa-Abstract.html.

Week 9: Midterm & Optimization and computation performance

Midterm will be held during the first half of the lecture. The second half will be on optimization and computation performance.

Topics:

- Gradient descent and stochastic gradient descent (SGD)
- Minibatch SGD
- Momentum
- Optimization algorithms: Adagrad, RMSProp, Adadelta, and Adam

Readings:

- Zhang et al, Chapter 12
- (Optional) Zhang et al, Chapter 13
- (Optional) Goodfellow et al, Chapter 8

Week 10: Natural language processing

Topics:

- Word embedding with skip-gram and continuous bag of words (CBOW) models (word2vec)
- Approximate Training with negative sampling and hierarchical softmax
- Word embedding with Global Vectors (GloVe)
- Bidirectional Encoder Representations form Transformers (BERT)

Readings:

- Zhang et al, Chapter 15
- (Optional) Zhang et al, Chapter 16

Week 11: Deep generative models

Topics:

- Boltzmann Machines
- Restricted Boltzmann Machines
- Deep belief networks
- Deep Boltzmann Machines
- Limited Boltzmann Machines and the DWave adiabatic quantum annealer
- Dynamic topology reconfiguration of Boltzmann Machines

Readings:

- Goodfellow et al, Chapter 20
- Liu, Jeremy, Ke-Thia Yao, and Federico Spedalieri. 2020. “Dynamic Topology Reconfiguration of Boltzmann Machines on Quantum Annealers.” *Entropy* 22 (11): 1202. https://doi.org/10.3390/e22111202.

Week 12: Generative Adversarial Networks

Topics:

- Generators and discriminators
- Training adversarial networks
- Deep convolutional generative adversarial networks.

Readings:

- Zhang et al, Chapter 18

Week 13: Graph neural networks

Topics:

- Graph convolutional networks
- Message passing
- Graph attentional networks

- Temporal graph networks
- Scaling graph neural networks

Readings:

- Hamilton, Chapter 3
- Hamilton, Chapter 5
- Wu, Zonghan, Shirui Pan, Fengwen Chen, Guodong Long, Chengqi Zhang, and Philip S. Yu. 2021. "A Comprehensive Survey on Graph Neural Networks." *IEEE Transactions on Neural Networks and Learning Systems* 32 (1): 4–24. https://doi.org/10.1109/TNNLS.2020.2978386.
- Rossi, Emanuele, Ben Chamberlain, Fabrizio Frasca, Davide Eynard, Federico Monti, and Michael Bronstein. 2020. "Temporal Graph Networks for Deep Learning on Dynamic Graphs." *ArXiv:2006.10637 [Cs, Stat]*, October. http://arxiv.org/abs/2006.10637.
- Wu, Felix, Tianyi Zhang, Amauri Holanda de Souza Jr., Christopher Fifty, Tao Yu, and Kilian Q. Weinberger. 2019. "Simplifying Graph Convolutional Networks." arXiv. https://doi.org/10.48550/arXiv.1902.07153.
- Chen, Ming, Zhewei Wei, Bolin Ding, Yaliang Li, Ye Yuan, Xiaoyong Du, and Ji-Rong Wen. 2020. "Scalable Graph Neural Networks via Bidirectional Propagation." In *Advances in Neural Information Processing Systems*, 33:14556–66. Curran Associates, Inc. https://proceedings.neurips.cc/paper/2020/hash/a7789ef88d599b8df86bbee632b2994d-Abstract.html.

Week 14: Knowledge graph embeddings

Topics:

- Link prediction and knowledge graph completion problem
- Translational distance models (TransE, TransH, TransR)
- Bilinear and tensor models (RESCAL, DistMult, TuckER)
- Graph neural network models (ConvE, ConvKB)

Readings:

- Hamilton, Chapter 4
- Wang, Q., Z. Mao, B. Wang, and L. Guo. 2017. "Knowledge Graph Embedding: A Survey of Approaches and Applications." *IEEE Transactions on Knowledge and Data Engineering* 29 (12): 2724–43. https://doi.org/10.1109/TKDE.2017.2754499.
- Lerer, Adam, Ledell Wu, Jiajun Shen, Timothee Lacroix, Luca Wehrstedt, Abhijit Bose, and Alex Peysakhovich. 2019. "PyTorch-BigGraph: A Large-Scale Graph Embedding System." *ArXiv:1903.12287 [Cs, Stat]*, April. http://arxiv.org/abs/1903.12287.
- Yu, Donghan, Yiming Yang, Ruohong Zhang, and Yuexin Wu. 2021. "Knowledge Embedding Based Graph Convolutional Network." In *Proceedings of the Web Conference 2021*, 1619–28. Ljubljana Slovenia: ACM. https://doi.org/10.1145/3442381.3449925.
- Lu, Yinquan, Haonan Lu, Guirong Fu, and Qun Liu. 2021. "KELM: Knowledge Enhanced Pre-Trained Language Representations with Message Passing on Hierarchical Relational Graphs." *ArXiv:2109.04223 [Cs]*, September. http://arxiv.org/abs/2109.04223.

Week 15: Project presentation and Poster Session

Students are expected to present their semester projects to the class, as one component of the semester project, in addition to the one-page project proposal, the project midterm report, and final project. In lieu of the final exam, the final project report is due on the day/time of the final.

Statement on Academic Conduct and Support Systems

Academic Integrity:
The University of Southern California is a learning community committed to developing successful scholars and researchers dedicated to the pursuit of knowledge and the dissemination of ideas. Academic misconduct, which includes any act of dishonesty in the production or submission of academic work, comprises the integrity of the person who commits the act and can impugn the perceived integrity of the entire university community. It stands in opposition to the university's mission to research, educate, and contribute productively to our community and the world.

All students are expected to submit assignments that represent their own original work, and that have been prepared specifically for the course or section for which they have been submitted. You may not submit work written by others or "recycle" work prepared for other courses without obtaining written permission from the instructor(s).

Other violations of academic integrity include, but are not limited to, cheating, plagiarism, fabrication (e.g., falsifying data), collusion, knowingly assisting others in acts of academic dishonesty, and any act that gains or is intended to gain an unfair academic advantage.

The impact of academic dishonesty is far-reaching and is considered a serious offense against the university. All incidences of academic misconduct will be reported to the Office of Academic Integrity and could result in outcomes such as failure on the assignment, failure in the course, suspension, or even expulsion from the university.

For more information about academic integrity see the student handbook or the Office of Academic Integrity's website, and university policies on Research and Scholarship Misconduct.

Please ask your instructor if you are unsure what constitutes unauthorized assistance on an exam or assignment, or what information requires citation and/or attribution.

Students and Disability Accommodations:
USC welcomes students with disabilities into all of the University's educational programs. The Office of Student Accessibility Services (OSAS) is responsible for the determination of appropriate accommodations for students who encounter disability-related barriers. Once a student has completed the OSAS process (registration, initial appointment, and submitted documentation) and accommodations are determined to be reasonable and appropriate, a Letter of Accommodation (LOA) will be available to generate for each course. The LOA must be given to each course instructor by the student and followed up with a discussion. This should be done as early in the semester as possible as accommodations are not retroactive. More information can be found at osas.usc.edu. You may contact OSAS at (213) 740-0776 or via email at osasfrontdesk@usc.edu.

Support Systems:

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

988 Suicide and Crisis Lifeline - 988 for both calls and text messages – 24/7 on call
The 988 Suicide and Crisis Lifeline (formerly known as the National Suicide Prevention Lifeline) provides free and confidential emotional support to people in suicidal crisis or emotional distress 24 hours a day, 7 days a week, across the United States. The Lifeline is comprised of a national network of over 200 local crisis centers, combining custom local care and resources with national standards and best practices. The new, shorter phone number makes it easier for people to remember and access mental health crisis services (though the previous 1 (800) 273-8255 number will continue to function indefinitely) and represents a continued commitment to those in crisis.

Relationship and Sexual Violence Prevention Services (RSVP) - (213) 740-9355(WELL) – 24/7 on call
Free and confidential therapy services, workshops, and training for situations related to gender- and power-based harm (including sexual assault, intimate partner violence, and stalking).

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

The Office of Student Accessibility Services (OSAS) - (213) 740-0776
OSAS ensures equal access for students with disabilities through providing academic accommodations and auxiliary aids in accordance with federal laws and university policy.

USC Campus Support and Intervention - (213) 740-0411
Assists students and families in resolving complex personal, financial, and academic issues adversely affecting their success as a student.

Diversity, Equity and Inclusion - (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: (213) 740-4321, HSC: (323) 442-1000 – 24/7 on call
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) 442-1200 – 24/7 on call

Non-emergency assistance or information.

Office of the Ombuds - *(213) 821-9556 (UPC) / (323-442-0382 (HSC)*
A safe and confidential place to share your USC-related issues with a University Ombuds who will work with you to explore options or paths to manage your concern.

Occupational Therapy Faculty Practice - *(323) 442-2850 or* otfp@med.usc.edu
Confidential Lifestyle Redesign services for USC students to support health promoting habits and routines that enhance quality of life and academic performance.