

Deep Learning in Medicine

BMSC-GA 4493, BMIN-GA 3007

Course Overview

The use of deep networks has revolutionized areas of image recognition, speech recognition, and natural language processing. Deep networks are also transforming the world of medicine by helping doctors to improve detection, diagnosis, treatment, and management of disease. Moreover, researchers begin to incorporate deep learning methods in areas such as drug development, diagnostic radiology and personalized medicine. In this course, we will focus on the deep learning approaches that are practical and currently used in various medical disciplines. In the labs, we will implement the approaches learned in the course using available multi-modal medical dataset.

Learning Objectives

Students will learn the most common deep learning methods emerging in medicine. Students will be able to differentiate various deep learning methods and choose the most appropriate ones for specific research problems.

General Information

Lectures - Tuesdays

- 10:00am-11:30am
- Lectures will be [live-streamed](#).

Lab Sessions - Thursdays

- 10:00am-11:30am
- Lab sessions will be given by teaching assistants and they will be [live-streamed](#).

Course Directors

Cem M. Deniz (cem.deniz@nyulangone.org)

Narges Razavian (narges.razavian@nyulangone.org)

Teaching Assistants

Haoxu Huang (hh2740@nyu.edu)

Tayyibah Khanam (tk2981@nyu.edu)

Naren Dhyani (naren.dhyani@nyu.edu)

Office Hours

Please send an email to TAs for an appointment.

Q&A Site

We will be using [Campuswire](#) for class discussion. The system is highly catered to getting you help fast and efficiently from classmates, the TA, and us. Rather than emailing questions to the teaching staff, we encourage you to post your questions on Campuswire.

Find our class page at: <https://campuswire.com/c/G30E4CA8B>

Lab and HW Github Repository

<https://github.com/nyumc-dl/BMSC-GA-4493-Spring2023>

Tools for Success

Python 3 is preferred, Jupyter Notebook, [PyTorch](#)

Course Materials

- **Book:** Ian Goodfellow, Yoshua Bengio and Aaron Courville (2016). *Deep Learning*, MIT Press. (Available in <http://www.deeplearningbook.org>)
- Weekly assigned articles and videos (*will be updated during semester*)

Prerequisites

Introduction to Programming (BMSC-GA 1358) or equivalent

Machine Learning (BMIN-GA 1004) or equivalent

Python Programming Language

Course Assessment

Homeworks (60%) + Final Project (40%: 5% Project Proposal + 5% Proposal Presentation + 25% Project Paper + 5% Paper Presentation)

Auditing

It is possible to audit the course for the NYU community.

Schedule (Tentative)

Day	Topic	Reading / Watching List (WiP)
01/24	Course Overview and Deep Learning Basics (CD, NR) (Lecture 0 - Course Logistics Lecture 1 - Intro to DL in Med)	• Deep Learning, Goodfellow et al, Chapter 2, Chapter 3 and Chapter 5. • LeCun Y, Bengio Y, Hinton G. Deep learning. <i>Nature</i>. 2015;521(7553):436-444. (pdf) • Mukherjee S. AI versus MD What happens when diagnosis is automated? <i>The New Yorker Annals of Medicine</i>. 2017. (link) • Deo RC, Machine Learning in Medicine, <i>Circulation</i>. 2015;132:1920-30 (pdf)

		Demis Hassabis et.al. "Neuroscience-Inspired Artificial Intelligence" <i>Neuron</i>. 2017.(link)
01/26	Lab 1: PyTorch and Packages Setup & Pytorch Tutorial (TK)	HW 1 Out
01/31	Deep Feedforward Networks (NR) (Lecture 2)	- Deep Learning, Goodfellow et al, Chapter 6. - Ravi D et al. Deep Learning for Health Informatics. <i>IEEE J Biomed Health Informatics</i>. 2017;21(1):4-21. (pdf) "Visualizing the Loss Landscape of Neural Nets" https://arxiv.org/pdf/1712.09913.pdf (Video: https://youtu.be/78vq6kgsTa8) - Playground: https://cs.stanford.edu/people/karpathy/convnetjs/demo/classify2d.html - More on basic calculus and differentiation (problems and solutions) https://www.math.ucdavis.edu/~kouba/ProblemsList.html
02/02	Lab 2: Deep Networks (HH)	
02/07	Convolutional Networks I (CD) (Lecture 3)	- Deep Learning, Goodfellow et al, Chapter 9. - Litjens G et al. A Survey on Deep Medical Image Analysis. 2017; <i>ArXiv: 1702.05747</i>. (pdf) - Lundervold et al. An overview of deep learning in medical imaging focusing on MRI. 2019; (link) - CNNs are explained very clearly: http://cs231n.github.io/convolutional-networks/
02/9	Lab 3: Pre-proposal meeting	
02/14	Convolutional Networks II (CD) (Lecture 4)	- Deep Learning, Goodfellow et al, Chapter 9. - Angermueller C et al. Deep Learning for Computational Biology. <i>Molecular Systems Biology</i> 2016; 12:878. (pdf) - Illustration of CNN architectures (link) 2D Visualization of a CNN (link)
02/16	Lab 4: CNNs	HW1 Due

- 02/21 **Recurrent and Recursive Nets I (NR) ([Lecture 5](#))**
- Deep Learning, Goodfellow et al, [Chapter 10](#).
 - Zachary C et al. Learning to Diagnose with LSTM Recurrent Neural Networks. 2016; *arXiv*: 1511.03677 ([pdf](#))
 - Mikolov, T., Sutskever, I., Chen, K., Corrado, G. S., & Dean, J. (2013). *Distributed representations of words and phrases and their compositionality*. Advances in neural information processing systems. ([pdf](#))
 - Benchmark: <https://gluebenchmark.com/leaderboard>
 - Sutskever, Ilya, Oriol Vinyals, and Quoc V. Le. "Sequence to sequence learning with neural networks." Advances in neural information processing systems. 2014. ([pdf](#))
 - Peters, M. E., Neumann, M., Iyyer, M., Gardner, M., Clark, C., Lee, K., & Zettlemoyer, L. (2018). Deep contextualized word representations. *arXiv preprint arXiv:1802.05365*. ([pdf](#))
- 02/23 Lab 5: CNN 2 (TK)
- 02/28 **Recurrent and Recursive Nets II: Attention Mechanism and Transformers (NR) ([Lecture 6](#))**
- Deep Learning, Goodfellow et al, [Chapter 10](#).
 - Original attention paper: <https://arxiv.org/pdf/1409.0473.pdf>
 - Transformers paper: ("Attention is all you need") <https://papers.nips.cc/paper/7181-attention-is-all-you-need.pdf>
 - OpenAI's first GPT model: https://s3-us-west-2.amazonaws.com/openai-assets/research-covers/language-unsupervised/language_understanding_paper.pdf
 - BERT (Bidirectional Encoder Representations from Transformers) <https://arxiv.org/abs/1810.04805>
- All below are optional but highly recommended:
- Huggingface NLP/Transformer course <https://huggingface.co/course/chapter1/1>

		- Transformer model implementations: https://huggingface.co/models - Andre Karpathy's "Let's build GPT" video https://www.youtube.com/watch?v=kCc8FmEb1nY (1.5 hour) - Attention model for image captioning: https://arxiv.org/abs/1502.03044
03/02	Lab 6: RNNs (ND)	
03/07	Student Project Proposal Presentations (one page papers are due at 10:00am)	
03/09	Feedback to Student Project Proposals	
03/14	No Class (Spring Break)	
03/16	No Class (Spring Break)	
03/21	Optimization of Deep Networks (CD) (Lecture 7)	- Deep Learning, Goodfellow et al, Chapter 4 and Chapter 8. - Mamoshina P et al. Applications of Deep Learning in Biomedicine. <i>Mol. Pharmaceutics</i> 2016; 13(5):1445-54. (pdf) - Sun R. Optimization for Deep Learning: Theory and Algorithms. <i>Arxiv 2019</i> (pdf) - Check visualizations of optimization algorithms: https://www.deeplearning.ai/ai-notes/optimization/
03/23	Lab 7: Optimization (ND)	HW2 Due (03/23) HW3 Out (03/23)
03/28	Regularization and Autoencoders (CD) (Lecture 8)	- Deep Learning, Goodfellow et al, Chapter 7 and Chapter 14. - Miotto R et al. Deep Patient: An Unsupervised Representation to Predict the Future of Patients

		from Electronic Health Records. <i>Scientific Reports</i> 6;26094. (pdf) • Ronneberger O et al. U-Net: Convolutional Networks for Biomedical Image Segmentation 2015; <i>arXiv:1505.0459</i> (pdf) • Srivastava N et al, Dropout: A Simple Way to Prevent Neural Networks from Overfitting. <i>Journal of Machine Learning Research</i> 2014; 15: 1929:1958. (pdf)
03/30	Lab 8: Autoencoders (TK)	
04/04	Practical Methodology (CD) (Lecture 9)	• Deep Learning, Goodfellow et al, Chapter 11 and Chapter 12. • Bengio Y. Practical Recommendations for Gradient-Based Training of Deep Architectures. In: Montavon G, Orr GB, Müller K-R, eds. <i>Neural Networks: Tricks of the Trade: Second Edition</i>. Berlin, Heidelberg: Springer Berlin Heidelberg; 2012:437-478 (pdf) • Smith L. N. Best Practices for Applying Deep Learning to Novel Applications; 2017 (pdf) • A Recipe for Training Neural Networks (link) • https://github.com/google-research/tuning_playbook
04/06	Lab 9 - Practical methods (ND)	
04/11	Transfer Learning (CD) (Lecture 10)	• Deep Learning, Goodfellow et al, Chapter 15. • He K, Chen X, Xie S, Li Y, Dollar P, Girshick R. Masked autoencoders are scalable vision learners. Arxiv 2021 (pdf) • Yosinski J, Clune J, Bengio Y, and Lipson H. How transferable are features in deep neural networks?. <i>NIPS</i>. 2014 (pdf) • Esteva A, Kuprel B, Novoa RA, et al. Dermatologist-level classification of skin cancer with deep neural networks. <i>Nature</i>. 2017;542(7639). (pdf) • Wang X, Peng T, Lu L, et al.

ChestX-Ray8: Hospital-Scale Chest X-Ray Database and Benchmarks on Weakly-Supervised Classification and Localization of Common Thorax Diseases. CVPR. 2017. ([pdf](#))

04/13	Lab 10: Transfer Learning (HH)	HW 3 Due (4/17)
04/18	Generative Adversarial Networks (NR) (Lecture 11)	• Deep Learning, Goodfellow et al, Chapter 15. • Goodfellow IJ et al. Generative Adversarial Networks. 2014; <i>arXiv:1406.2661</i> (pdf) <i>Binge Watch</i> • NIPS 2016 Workshop on Adversarial Training - Ian Goodfellow - Introduction to GANs (link)
04/20	Lab 11: GANs (HH)	
04/25	Advanced Topics (NR) (Lecture 12)	• Self-supervised learning • Graph Neural Networks
04/27	Lab 12: Self-supervised Learning (HH)	
05/02	Student Project Presentations I	
05/04	Student Project Presentations II	
05/06	Student Project Paper Deadline (5:00 pm)	

Homeworks

There will be three homework assignments, each will be graded for 20% of your final grade.

HW1 (due 02/16 11:59 pm): Back-Propagation and MLP

HW2 (due 03/23 11:59 pm): CNN

HW3 (due 04/17 11:59 pm): RNN

HW submission instructions: Students should submit a **zipped** folder named **netid_hwx** where x is the hw number . The submission should consist of the **jupyter notebook with all the plots and expected outputs clearly visible** in it. The zipped folder should also contain the data files. We should be able to **run your ipynb without making directory changes**. Not following the protocol might lead to deduction of scores.

Student Projects

Students are expected to work on a project related to the topics presented in the lectures.

Teams

You may work in teams of two on the final project. Teams must be finalized by the project proposal deadline (03/08). Each team member must build and test at least one deep learning model.

Topics

Students are free to choose their project topics related to medicine. Example research topics are:

- https://grand-challenge.org/All_Challenges
- Challenges or data from <https://www.kaggle.com/>
- Protein Sequence and Structures (<https://github.com/aqlaboratory/proteinnet>)
- Cancer Imaging Archive TCIA (<http://www.cancerimagingarchive.net/>)
- TCGA portal for genomics/pathology/proteomics (<https://portal.gdc.cancer.gov/>)
- Histopathology (Lymph nodes -> metastatic) (<https://github.com/basveeling/pcam>)
- Tadpole data for Alzheimer's (<https://tadpole.grand-challenge.org/Data/>)
- dreamchallenges.org (usually genomics)
- MURA dataset
- Breast Mammography
- [NIH Chest X-ray disease identification](#), [MIMIC CXR chest x-ray data](#) (371,920 chest x-rays associated with 227,943 imaging studies), [CheXpert](#) (Stanford chest x-ray data 224,316 chest radiographs of 65,240 patients)
- [MIMIC III Critical Care Data](#) (58,000 hospital admissions for 38,645 adults and 7,875 neonates. The data spans June 2001 - October 2012) * This dataset is massive/complex. If you want to work on this you need to email us by the end of the ****first**** week of the semester.
- MIMIC IV Data (Includes Chest XRay images **MIMIC-IV-CXR**, paired with ED department data **MIMIC-IV-ED**, notes **MIMIC-IV-Note**, etc) over 65K Patients: <https://mimic.mit.edu/>
- [Cancer data from the National Cancer Institute's CPTAC program](#) <https://pypi.org/project/cptac/>
- <https://www.kaggle.com/c/hpa-single-cell-image-classification>
- <https://fastmri.org/>

Pre-proposal meeting

Before your proposal deadline, you must discuss your ideas and especially your dataset with the instructors and TAs.

Proposal

You must submit your one page proposal and make a presentation to receive useful feedback on the project.

For the proposal use the following Overleaf template: <https://www.overleaf.com/read/jirkissjxtb>

Upload and Submit PDF file and your presentation in newclasses.nyu.edu

Paper and Presentation

The final report should include the description of the task, models, experiments and conclusion. It can be up to **6 pages long** excluding unlimited pages reserved for references. You **MUST** follow the formatting for JMLR articles (<http://www.jmlr.org/format/format.html>). You can use Overleaf template from the following link: <https://www.overleaf.com/latex/templates/template-for-journal-of-machine-learning-research-jmlr-with-jmlr2e-dot-sty/vjcpvhvztrjn>

We use the following rubric to grade project papers:

<https://docs.google.com/spreadsheets/d/186-zpvGhCyQrkkPxzzj1MrSxVQ5qWjl77pQ4beB7fs/edit?usp=sharing>

In the final week of the semester, you will present your findings in the class. Both of your paper and presentation will affect your grade on the student project.

Contribution Statements

The final report must state the contributions of each team member. Each team member should indicate which model they implemented. Team projects which fail to include this will receive a 5% grade deduction.

General Policies

Late/missed work

You must adhere to the due dates for all required submissions. If you miss a deadline, then you will not get credit for that assignment/post.

Bonus Questions in Homeworks

It may be possible that we will add bonus questions in the homework. These bonus scores will **not** carry over to the final score, and are only there to compensate within that homework. Only in special circumstances when these bonus questions would change a student's status from fail to pass, we will consider carrying the bonus scores over to adjust the final score.

Incomplete

No "Incomplete" will be assigned for this course unless we are at the very end of the course and you have an emergency.

Responding to Messages

We use Campuswire for all course related questions/discussions related to the course. If you have highly private/sensitive issues and prefer to communicate via email, please keep in mind that: we will check emails daily during the week, and we will respond to you within 48 hours.

Announcements

We will make announcements throughout the semester by e-mail.

Make sure that your email address is updated; otherwise you may miss important emails from us.

Safeguards

Always backup your work in a safe place (electronic file with a backup is recommended) and make a hard copy. Do not wait for the last minute to do your work. Allow time for deadlines.

Plagiarism

Plagiarism, the presentation of someone else's words or ideas as your own, is a serious offense and will not be tolerated in this class. The first time you plagiarize someone else's work, you will receive a zero for that assignment. The second time you plagiarize, you will fail the course with a notation of academic dishonesty on your official record.