

SIADS 642 Introduction to Deep Learning

Fall 2020 Syllabus

Course Overview and Prerequisites

This course introduces the basic concepts of Neural Networks and Deep Learning. Students will learn the basic model types used in Deep Learning and their suitability for various data domains such as text, images, and videos. By the end of this course, students will be able to extract patterns from real-world datasets by running several classes of deep learning methods using existing code via standard API calls.

The course prerequisites are SIADS 502: Math Methods for Data Science, SIADS 505: Data Manipulation, and SIADS 542: Supervised Learning.

Instructor and Course Assistants

- Instructor: Paramveer Dhillon (dhillonp@umich.edu)
- Course Assistant: Yumou Wei (yumouwei@umich.edu)

Communication Expectations

Contacting instructor and course assistant: Please only use the course channel in Slack. Please ask all questions in public if possible so others can learn from your question. If you need to ask a private question please direct message both the instructor and the course assistant.

Email response time: 48 hours

Slack response time: 48 hours

Office hours: see the [Google Calendar](#).

Textbook (Optional)

[Deep Learning](#) by *Goodfellow, Bengio, and Courville*. Available at the link for free. Copyright 2016 MIT Press.

Learning Outcomes

1. Learn how to modify state-of-the-art deep learning architectures for a new dataset/task.
2. Know the basic model types used in deep learning, e.g., Convolutional Neural Networks (CNNs), Recurrent Neural Networks (RNNs), and Generative Adversarial Networks (GANs).

3. Understand the basic concepts of neural networks and deep learning methods.
4. Know the suitability of specific deep learning methods to various real world data domains such as the ones arising from text, images, and videos.
5. Be able to re-train and tune hyperparameters of several classes of deep learning methods, in particular CNNs, RNNs, and GANs, on real-world datasets.
6. Be able to extract patterns from complex real world image and text datasets by using deep learning methods.

Course Schedule

This course **begins on Monday, October 26, 2020** and **ends on Sunday, November 22, 2020**.

Assignments will be due as noted in the Coursera platform (roughly weekly).

Week 1:

Introduction to Deep Learning.

Topics Covered: a) What is deep learning? What are Neural Networks?, b) The basic building blocks of deep learning, c) Training of Neural Networks.

Learning outcomes: Know the basic concepts and terminology of neural networks and deep learning. Understand the issues involved in training deep learning models and the various "tricks" commonly used to get good model performance. Describe how deep learning models are related to standard machine learning models.

Week 2:

Convolutional Neural Networks (CNNs)

Topics Covered: a) What are CNNs and why do we need them?, b) The basic building blocks of CNNs. CNN architectural details, c) Study of some famous CNN architectures.

Learning outcomes: Know and understand the details of CNNs including their architectural details and training. Be aware of the different CNN architectures used in literature. Describe how and which real-world problems are best modeled using CNNs.

Week 3:

Recurrent Neural Networks (RNNs)

Topics Covered: a) What are RNNs and why do we need them?, b) The basic building blocks of RNNs and other architectural details, c) Study of some special RNN architectures: GRU, LSTM, Seq2Seq models, Attention mechanism.

Learning outcomes: Know and understand the details of RNNs including their architectural details and training. Understand the working details of popular RNN extensions such as GRUs, LSTMs. Describe how and which real-world problems are best modeled using

RNNs. Describe the benefits of advanced RNN models such as Seq2Seq and Attention Mechanism.

Week 4:

Adversarial Learning Models

Topics Covered: a) What are generative models? How are they different from discriminative models?, b) Architectural and training details of Generative Adversarial Networks (GANs), c) Adversarial attacks on Neural Networks.

Learning outcomes: Understand the principles of Adversarial Learning Models. Know the architectural and training details of GANs. Describe how Adversarial Attacks can cripple the neural networks. Describe how deep generative model architectures are principally different from CNNs and RNNs described in previous weeks.

Grading

Course Assignment	Percentage of Final Grade
Week 1 Multiple-choice quiz	7%
Week 1 Jupyter Notebook Assignment	20%
Week 2 Multiple-choice quiz	7%
Week 2 Jupyter Notebook Assignment	20%
Week 3 Multiple-choice quiz	6%
Week 3 Jupyter Notebook Assignment	20%
Week 4 Jupyter Notebook Assignment	20%

Letter Grades, Course Grades, and Late Submission Policy

Refer to the [MADS Assignment Submission and Grading Policies](#) section of the UMSI Student Handbook (access to Student Orientation course required)

Late Submission Policy: 1 day late = 15% reduction; 2 days late = 30%; 3 days late = 50%; 4 days late = 0 points. Extenuating circumstances will be considered (please reach out to the instructor as soon as possible).

Percentage grades will be converted to letter grades using the following formula:

Total Assignment Scores	Letter Equivalent
100	A+
95-99	A
90-95	A-
85-90	B+
80-85	B
75-80	B-
70-75	C+
65-70	C
60-65	C-
55-60	D+
50-55	D
<50	E

Accommodations

Refer to the [Accommodations for Students with Disabilities](#) section of the UMSI Student Handbook (access to the Student Orientation course required).

Use the [Student Intake Form](#) to begin the process of working with the University's Office of Services for Students with Disabilities.

Accessibility

[Screen reader configuration for Jupyter Notebook Content](#)

Help Desk(s): How to get Help

- Degree program questions or general help - umsimadshelp@umich.edu
- Coursera's Technical Support (24/7) - <https://learner.coursera.help/>

Library Access

Refer to the [U-M Library's information sheet](#) on accessing library resources from off-campus. For more information regarding library support services, please refer to the [U-M Library Resources](#) section of the UMSI Student Handbook (access to the Student Orientation course required).

Student Mental Health

Refer to the University's [Resources for Stress and Mental Health website](#) for a listing of resources for students.

Student Services

Refer to the [Introduction to UMSI Student Life](#) section of the UMSI Student Handbook (access to the Student Orientation course required).

Office Hours

Paramveer Dhillon - Mondays 10-11am ET

Yumou Wei - 9-10pm ET