

Information Technology
Course Number: 22:544:635
Course Title: Neural Networks and Deep Learning

COURSE DESCRIPTION

This course introduces modern techniques of neural networks and deep learning, which have revolutionized machine learning and artificial intelligence practice to graduate students. The course heavily relies on software and libraries for deep learning, including Tensorflow, Keras, and similar tools, and students are required to conduct extensive projects. An end of term team project and presentation in class is also required. Students are exposed to extremely large data with a complex feature set. Pattern and image recognition, speech and sentiment analysis, and generating new images, texts, painting, and sounds from a given set of are among some application areas explored. A rigorous review of basic statistical foundations of neural networks, the maximum likelihood technique is presented. Basic optimization techniques used in neural networks are reviewed, including gradient descent, stochastic gradient descent, Momentum techniques such as RMSPROP, and Nesterov's acceleration are covered, and ADAM and NADAM algorithms are reviewed. Various regularization techniques, from conventional L2 and L1 to dropout and early stopping techniques, are also reviewed. An in-depth analysis of the backpropagation algorithm is presented. Various models of neural networks and their applications are covered: sequential dense feedforward networks, one and twodimensional convolutional networks, recurrent networks, Boltzman machines, autoencoders, variational autoencoders, and Generative Adversarial Networks. Both supervised learning in regression and classification, and unsupervised techniques such as feature extraction and generative models are covered.

COURSE MATERIALS

1. I. Goodfellow, Y. Bengio and A. Courville, *Deep Learning*, MIT Press 2016. (Available online at <http://www.deeplearningbook.org>)
2. Aurélien Géron, "Hands-on Machine Learning with Scikit-Learn, Keras & TensorFlow" O'Reilly publishers
3. F. Chollet, *Deep Learning with Python*, second edition Manning, 2021.

The first book is for theoretical understanding, the second and third for learning how to use the Keras software along with Tensorflow.

Additional code scripts and lecture notes may be added at the instructor's discretion.

LEARNING GOALS AND OBJECTIVES

- Mastery and in-depth understanding of neural network architectures and deep learning: among them Dense, convolutional, recurrent, autoencoders, variational autoencoders, generative adversarial networks, and general Boltzman Machines
 - Mastery and in-depth understanding of inner workings of neural networks, in particular, various hyperparameters, activation functions choices, optimization algorithms, various regularization techniques, and various startup approaches
 - Students will become familiar with software development and implementation of neural network techniques, and will learn how to look for the most suitable libraries for their tasks
 - Students will learn to access very large data sets from various sites, learn how to model proper neural architecture for their particular problem, and learn how to leverage available resource, for example Google Collab, Amazon AWS, or Microsoft Azure to perform computing entirely online and using the most sophisticated and advanced hardware available in such sites, they will learn to access vast and real-world data sets from sites such as Kaggle
 - Students will learn how to collaborate in a team project for nontrivial, exciting, relevant and practical learning projects. They will learn how to showcase their work on sites such as Github.
-

PREREQUISITES

- A good knowledge of undergraduate level linear algebra and calculus. Also, a good knowledge of basic, undergraduate level probability and statistics.
 - 22:544:531 or similar course on a general machine learning.
 - Students must be familiar with the Python programming language.
-

ACADEMIC INTEGRITY

I do NOT tolerate cheating. Students are responsible for understanding the RU Academic Integrity Policy (<http://academicintegrity.rutgers.edu/>)

I will strongly enforce this Policy and pursue *all* violations. On all examinations and assignments, students must sign the RU Honor Pledge, which states, “On my honor, I have neither received nor given any unauthorized assistance on this examination or assignment.” I will screen all written assignments through *SafeAssign* or *Turnitin*, plagiarism detection services that compare the work against a large database of past work. Don’t let cheating destroy your hard-earned opportunity to learn. See business.rutgers.edu/ai for more details.

ATTENDANCE AND PREPARATION POLICY

- Attendance is required, and students are responsible for all information conveyed in classes. While some information may also be shared on the course web site, students should not assume that all pertinent information will be found there. If for some reason you cannot attend the lecture let the instructor know through Canvas message tool (no e-mail, please.)
- For weather emergencies, consult the campus home page. If the campus is open, class will be held.

- Expect me to arrive on time for each class session. I expect the same of you. [If you are going to be tardy, then please send me an email to let me know in advance.]
 - Expect me to remain for the entirety of each class session. I expect the same of you. [If you are going to leave early, then please send me an email to let me know in advance.]
 - Expect me to prepare properly for each class session. I expect the same of you. Complete all background reading and assignments. You cannot learn if you are not prepared. The minimum expectation is that for each [X]-hour class session, you have prepared by studying for at least twice as many hours.
 - Expect me to participate fully in each class session. I expect the same of you. Stay focused and involved. You cannot learn if you are not paying attention.
-

CLASSROOM CONDUCT

- Please silence your cell phones during the lecture time.
 - No side conversations, sleeping, phone conversations and texting during lectures please.
-

EXAM DATES AND POLICIES

There will be no exams in the course. Grading is entirely based on class projects, and end term project, and class participation.

GRADING POLICY

Course grades are determined as follows:

40% Class projects
50% Final project
10% Class participation

- There will be 4-5 homework projects once around every two weeks. Homeworks will include programming exercises as well as conceptual questions.
 - Extra credit: Extra credit is available for students with excellent class participation.
 - Grade posting: The grades will be posted on the Canvas.
-

COURSE SCHEDULE

- Week 1. An overview of neural networks as a machine learning tools, and introduction to Keras layer building system
- Week 2. A review of machine learning basics: Bayes decision rule, supervised learning and regression and classification, forms of data, loss functions, risk, empirical risk. Review of the maximum likelihood paradigm
- Week 3. An introduction to sequential, dense neural networks and their applications, review of various activation functions: linear, cross-entropy, softmax, Relu, and related, case studies on the MNIST handwritten digit data, Keras and cloud computing platforms for neural networks are introduced
- Week 4. Continuation of Week 3
- Week 5. A review of optimization techniques: Notion of gradient, subgradient and Hessian, convex and nonconvex functions, gradient descent, stochastic gradient descent, momentum methods such as RMSPROP and Nesterov's acceleration, ADAM and NADAM, description and analysis of the back propagation algorithm. Introduction to Tensorflow and how complex functions, along with their gradients and Hessians, can be represented as Tensorflow computation graph, and evaluated automatically and efficiently
- Week 6. Continuation of Week 5
- Week 7. Regularization techniques and their impact; L2 and L1 regularization, Dropout, dataset augmentation, early stopping
- Week 8. Convolution as a method of feature extraction, 2D convolutional neural network for image recognition and classification. Use of pre-trained networks, Image augmentation techniques (rotation, reflection, translation, zoom-in/zoom-out), pooling techniques 1D convolutional networks and their applications in time series data, word prediction, sentiment prediction.
- Week 9. Continuation of Week 8
- Week 10. Continuation of Week 9
- Week 11. Recursive functions and their optimization, representation of recursive functions by Recurrent networks, simple and unfolded representation of recurrent neural networks, equivalence of recurrent networks and Turing Machines, LSTM and GRU units for long term memory, word embedding, case studies in text understanding, word generation, sentiment detection, classifying text according to subject,
- Week 12. Continuation of Week 11 plus Attention and Transformation networks GPT-3
- Week 13. Review of Principal Component Analysis (PCA) and its probabilistic version, introduction to autoencoders and encoder decoder networks, relation to PCA
- Week 14. Generative models, Boltzman machine, variational autoencoders for generation, Generative Adversarial Networks (GAN), constructing probability distribution of certain very high

dimensional data, applications of generative models, case studies of generating handwritten digits, pictures of animals, music, writings in certain style

SUPPORT SERVICES

If you need accommodation for a *disability*, obtain a Letter of Accommodation from the Office of Disability Services. The Office of Disability Services at Rutgers, The State University of New Jersey, provides student-centered and student-inclusive programming in compliance with the Americans with Disabilities Act of 1990, the Americans with Disabilities Act Amendments of 2008, Section 504 of the Rehabilitation Act of 1973, Section 508 of the Rehabilitation Act of 1998, and the New Jersey Law Against Discrimination. More information can be found at ods.rutgers.edu.

[Rutgers University-New Brunswick ODS phone (848)445-6800 or email dsoffice@echo.rutgers.edu]

[Rutgers University-Newark ODS phone (973)353-5375 or email ods@newark.rutgers.edu]

If you are *pregnant*, the Office of Title IX and ADA Compliance is available to assist with any concerns or potential accommodations related to pregnancy.

[Rutgers University-New Brunswick Title IX Coordinator phone (848)932-8200 or email jackie.moran@rutgers.edu]

[Rutgers University-Newark Office of Title IX and ADA Compliance phone (973)353-1906 or email TitleIX@newark.rutgers.edu]

If you seek *religious accommodations*, the Office of the Dean of Students is available to verify absences for religious observance, as needed.

[Rutgers University-New Brunswick Dean of Students phone (848)932-2300 or email deanofstudents@echo.rutgers.edu]

[Rutgers University-Newark Dean of Students phone (973)353-5063 or email DeanofStudents@newark.rutgers.edu]

If you have experienced any form of *gender or sex-based discrimination or harassment*, including sexual assault, sexual harassment, relationship violence, or stalking, the Office for Violence Prevention and Victim Assistance provides help and support. More information can be found at <http://vpva.rutgers.edu/>.

[Rutgers University-New Brunswick incident report link: <http://studentconduct.rutgers.edu/concern/>. You may contact the Office for Violence Prevention and Victim Assistance at (848)932-1181]

[Rutgers University-Newark incident report link:

https://cm.maxient.com/reportingform.php?RutgersUniv&layout_id=7 . You may also contact the Office of Title IX and ADA Compliance at (973)353-1906 or email at TitleIX@newark.rutgers.edu. If you wish to speak with a staff member who is confidential and does **not** have a reporting responsibility, you may contact the Office for Violence Prevention and Victim Assistance at (973)353-1918 or email run.vpva@rutgers.edu]

If students who have experienced a temporary condition or injury that is adversely affecting their ability to fully participate, you should submit a request via <https://temporaryconditions.rutgers.edu> .

If you are a military **veteran** or are on active military duty, you can obtain support through the Office of Veteran and Military Programs and Services. <http://veterans.rutgers.edu/>

If you are in need of **mental health** services, please use our readily available services.

[Rutgers University-Newark Counseling Center: <http://counseling.newark.rutgers.edu/>]

[Rutgers Counseling and Psychological Services–New Brunswick: <http://rhscaps.rutgers.edu/>]

If you are in need of **physical health** services, please use our readily available services.

[Rutgers Health Services – Newark: <http://health.newark.rutgers.edu/>]

[Rutgers Health Services – New Brunswick: <http://health.rutgers.edu/>]

If you are in need of **legal** services, please use our readily available services: <http://rusls.rutgers.edu/>

Students experiencing difficulty in courses due to **English as a second language (ESL)** should contact the Program in American Language Studies for supports.

[Rutgers–Newark: PALS@newark.rutgers.edu]

[Rutgers–New Brunswick: eslpals@english.rutgers.edu]

If you are in need of additional **academic assistance**, please use our readily available services.

[Rutgers University-Newark Learning Center: <http://www.ncas.rutgers.edu/rlc>

[Rutgers University-Newark Writing Center: <http://www.ncas.rutgers.edu/writingcenter>]

[Rutgers University-New Brunswick Learning Center: <https://rlc.rutgers.edu/>]

[Optional items that many faculty include:

- Students must sign, date, and return a statement declaring that they understand the RU Academic Integrity Policy.

- Students must sign, date, and return a statement declaring that they understand this syllabus.]