Neural Machine Learning I

COMP / ELEC / STAT 502, Fall 2024

Elective course, 3 credits (3 contact hours per week)

Course home page: https://www.ece.rice.edu/~erzsebet/ANNcourse.html
(This website is also linked from Canvas.)
Students are responsible for being familiar with this syllabus, and with the contents of both the above and the Canvas website and follow the postings as the course proceeds.

Class meets: TR 10:50am – 12:05pm in MXF 252
In-person, in-class instruction is planned, except for possible Rice-mandated periods of remote instruction*
Please plan to attend synchronously even if we may conduct classes remotely via Zoom for a short period. Classes will not be recorded.

Instructor: Erzsébet Merényi
email: erzsebet@rice.edu
Office/Phone: MXF 229, 713-348-3595
Office hour: by appointment Preferred window: Tuesday 4:00 - 5:30pm (subject to change)
Teaching Assistant and contact information: TBA
TA advising: by appointment. Preferred window TBA.

The information contained in the course syllabus, other than the absence policies, may be subject to change with reasonable advance notice as deemed appropriate by the instructor.

*In particular, if / when Covid-19 should force substantial changes in safety measures we may switch to remote instruction via Zoom for any part of the semester as deemed necessary by University leadership. In that case changes will be posted in a timely manner at the above course web site and on the Canvas home page, as well as via Canvas announcements. Any related logistic issues will be posted / discussed at the same time. Under Covid-19 safety measures, while in class or visiting me in my office I expect everyone to strictly adhere to Rice University's mask requirements (as applicable) or other guidelines for minimizing the spread of the Covid-19 virus.

Short course description

Review of major Artificial Neural Network paradigms. Analytical discussion of supervised and unsupervised learning. Emphasis on state-of-the-art Hebbian (biologically most plausible) learning paradigms and their relation to information theoretical methods. Applications to data analysis such as pattern recognition, clustering (information discovery), classification, non-linear

PCA, independent component analysis, with examples and exercises from image and signal processing and other areas.

Brief Sample List of Topics

Introduction: What is an Artificial Neural Network (ANN), defining characteristics
Review of Basics: from Calculus, Information Theory and Statistics
General Machine Learning concepts: supervised and unsupervised learning; training and test data, cross-validation
Associative Memory: recall from partial and noisy samples, stability
Supervised Lerning: Simple and Multi-Layer Perceptrons
Evaluation of Learning and Generalization Performance
Unsupervised Learning: Hebbian learning, PCA nets, Competitive Learning (SOMs, LVQs)
Recurrent nets: Hopfield and Boltzmann Machines, Simulated Annealing
Speeding Up Supervised Learning, and Structure Optimization
Objective Functions for ANN Learning: information theoretical functions
Identifying Independent Sources (Blind Source Separation) with ANNs

A more detailed sample course outline is posted at
https://www.ece.rice.edu/~erzsebet/ANNcourse/Sample-Syllabus_COMP-ELEC-STAT-502.pdf

Exercises, homeworks, exam and course project will involve programming in Matlab (preferred), or C, R (Fortran also acceptable). Code handouts will be in Matlab.

Detailed Course Schedule

A detailed schedule of class topics will come on-line in a timely manner at the course web site https://www.ece.rice.edu/~erzsebet/ANNcourse.html under Course Schedule along with reading assignments from text book and lecture notes. I will also indicate there when home work assignments are posted in Canvas, along with due dates; post the dates of upcoming tests, and all relevant logistics. The materials (such as lecture notes and home work assignments) indicated in the Course Schedule will be downloadable from Canvas, and you will submit assignments electronically through Canvas.

Pre-requisite

- Students are generally expected to handle mathematics / probability / statistics on 300 level, 400 level is better.

Details on desired background

- No previous knowledge of Artificial Neural Nets is assumed.
- Linear algebra (such as in CAAM 335, MATH 355, ELEC 301, ELEC 303, or equivalent), and multivariate calculus (such as in CAAM 336, MATH 321, CAAM 501, or equivalent);

- Probability and statistics (such as in STAT 310, or ELEC/STAT 331 or equivalent);
- Simple information theoretical notions (ELEC 241 or equivalent); These will be briefly reviewed in the course.
- If you have taken (ELEC 531 and (CAAM 583 / ELEC 533 / STAT 583)) or (STAT 615 and STAT 518) you automatically qualify.

Additionally, courses such as the following (or similar) are an advantage: ELEC/COMP 440 or ELEC 535; ELEC 478, 483, CAAM 416, 453; ELEC / CAAM / MECH 508; STAT 413 / 613

Course Materials

The course will be based on Lecture Notes, scheduled as described at the course site https://www.ece.rice.edu/~erzsebet/ANNcourse.html under Course Schedule .

Recommended reading will be selected parts of the following books. All are available at www.amazon.com, or at the Fondren Library at Rice. You can also borrow my copies for short periods of times.

On neural networks

- Frederick Ham and Ivica Kostanic: Principles of Neurocomputing for Science & Engineering. McGraw-Hill, 2001.
- Simon Haykin: Neural Networks. A Comprehensive Foundation. McMillan, New Jersey, 1999. (2nd Edition)
- Christopher M. Bishop: Neural Networks for Pattern Recognition. Oxford University Press, 1995.

On machine learning (neural and other)

- Christopher M. Bishop: Pattern Recognition and Machine Learning. Springer, 2007.

Further suggested reading is listed at the course web site under Course Materials, and occasional articles will be suggested in the Course Schedule and provided in Canvas.

Objectives of the Course

1. Student understanding of concepts and mastery of the main types of neural machine learning models and methods, and their applications to data mining, clustering, classification and regression.

2. Student mastery of both theoretical approaches and application of neural network simulations for solving data analysis problems. Understanding and practicing correct methodology for good neural network learning and generalization, and assessment of the quality of the learning (the generalization capability of trained networks).

3. Student competence in communicating the methods and results of data analyses with neural networks.

Assignments, Grading Policies and Other Logistic Requirements

Grades will be made up of the following components, with approximate weights as shown:

25% - Homework assignments
7.5% - Quiz 1
7.5% - Quiz 2
25% - Exam 1
30% - Course Project
5% - Points earned in bonus & other challenges (aka Pizza Points)

This break down is ***subject to change depending on resources that will be available to the instructor for grading vs the class size***. Parts of points 1. – 4. below may change accordingly.

As of January 8, 2019, the grading capacity became insufficient for full grading of individual homeworks. Therefore, students do homework in groups. The groups, formed by self-organization 😊, are permanent throughout the semester. Maximum group sizes will be determined by the instructor. Students in a group work together, agree on, and produce a homework solution and write-up, submit it as a group. Each student in a group will receive the same grade points for the homework unless they state contributions in a proportion different from equal percentage. In each homework, they will state the names of the group members, and the percentages of their contributions if unequal.

The homework grouping does not affect quizzes and exams, which will be done individually and will be graded individually. Class projects (see below), however, will be done in groups. At project time students will be allowed to form different groups from the homework groups if they want to.

1. Homework assignments

There will be approximately four or five homeworks in the first half of the semester. There will be only two homeworks in the second half of the semester, to allow time for the Course Project. The assignments themselves will be posted in Canvas. Homeworks will be due, by default, before the beginning of class on the due date as specified in Canvas. Due dates will also be shown in the Course Schedule in a timely manner. Solutions from previous years are off-limit. Partial grading may be applied (i.e., grading a randomly selected subset of the homework problems) depending on the number of students and grading support for the course.

You(r group) will be required to upload an electronic copy (a pdf file) to Canvas. You will not be required to type everything but if you have parts produced non-electronically you will need to scan your hard copy into a pdf file with a reasonably high-quality scanner.
A summary of required formats, file names, and other requirements are posted here in a single file, for your convenience. I will refer to this file as “Logistic Details”. Please be sure to thoroughly read and follow. You may save the file for your handy reference; or you can always

access it at https://www.ece.rice.edu/~erzsebet/ANNcourse/NML502-LogisticsDetails.pdf

1.1 Late homework policy

Homework will be due on the due date and time posted with the Assignment in Canvas. If the deadline indicated in the Course Schedule should be different, then the deadline posted in Canvas applies. After the due date, but before the late deadline (the "accept until" date in Canvas), homeworks can be turned in for 50% credit. After sample solutions are handed out, 0% credit will be issued. Should any circumstance prevent electronic submission to Canvas at any time, please *contact me before the assignment is due*, and copy the TA(s). We will work with you to solve the problem. However, please let me and the TA(s) know immediately (by phone if necessary) of any problem accessing Canvas.

2. Tests

Quiz 1 and Quiz 2: approx. 45 minutes, closed notes/books/phones/computers, planned as electronic "in-class" style quizzes. (An electronic device may be used only for writing and submitting the quiz answers.) The exact implementation (multiple-choice form, questions requiring short, typed answers, or combination) will be announced well ahead of time.

Exam 1 (midterm): Exam 1 will be take-home, open book and open notes but using restricted time and other resources (such as internet). For simulations, you will be allowed to use codes you developed in home works. You will be required to adhere to the Rice honor code. The exact requirements and details will be given with the exam. Test solutions from previous years are strictly off-limit.

The exact dates of these tests will be posted in the Course Schedule in a timely manner, as well as announced in class approximately two weeks ahead of time. The anticipated *approximate* dates are mid October (Quiz 1), late October - early November (Midterm), and mid November (Quiz 2). For all tests you will be required to adhere to the Rice Honor Pledge.

3. Course Project

A Project will be done over an approximately four-week period in the second half of the semester. Students will form groups by self-organization 😊. The project groups can be the same or different from the homework groups. Each group will submit a project proposal by approximately October 29 (exact date TBD, will be announced in class and will be posted in the Course Schedule). It is anticipated that by that time major ANN paradigms will have been discussed in class and students will be able to select a project topic, either devised on their own or selected from a list provided by me. I will review and approve each proposal (or discuss necessary modifications) before students start executing plans. Details, requirements and the characteristics of a successful Project are described here.

4. Bonus & other challenges will consist of competitions added to assignments, short in-class surveys (one or two questions in a Google form), and questions I post in Piazza, and can earn you "Pizza Points" that count toward your final grade. While these challenges are low-risk

individually (each counts very little toward your grade) you will benefit from instantaneous feedback and/or class discussion as you are accumulating them.

Missed assignments If you must miss (or be late with) a homework, or a test due to an extraordinary circumstance please notify me and the TA(s) *as much ahead of time* as possible, and **make arrangements with me** for completing the missed assignment. If, in extreme emergency, you are unable to provide advance notice, please let me know as soon as possible afterwards, and I will work with you on a solution accordingly. Please do not ask the TA(s) for extensions, make-up arrangements, and other exceptions. Address all such requests directly to me (but copy the TA(s)).

Other Requirements

In addition to file formats and file name conventions, the "Logistic Details" file (above) contains a summary of other logistic requirements such as the procedure that you must follow for disputing / appealing grade points; and email communication with me and TA(s) and graders. Please be sure to know and follow those requirements.

Class Attendance and Absence Policy

Students are strongly advised to attend classes and participate in class activities. Students who must miss a class or assignment because of unavoidable circumstances should consult with the instructor as much in advance as possible so that arrangements may be made to help you keep up with the material.

Expectations Regarding Honor Code, Collaboration, and Citation

In preparing homework, students are encouraged to consult freely any material and anyone. However, each individual (or each homework group) will write and turn in their own solution, which they are expected to understand and be able to reproduce on their own. If the solution is the result of group work, each group will write the names of the collaborators on their solution. For tests, students will be required to work individually, and adhere to the Rice Honor Code. Solutions from previous years are off-limit. In all work, students are expected to be scrupulous about proper citation of sources (where applicable), as required both as a matter of integrity and formally as a part of the Rice Honor Code.

University Disability Accommodation Policy

The University seeks to foster an environment of broad access and feasibly equal opportunity to education. The Office of Disability Support Services (DSS; Allen Center, Room 111; 713-348-584; adarice@rice.edu) supports and implements federal guidelines under the Rehabilitation Act of 1973 and the Americans with Disabilities Act. Students with

documented disabilities requiring accommodation under Rice's established policies should consult DSS and the instructor; all such consultations and accommodations will be held confidential to the extent feasible.

Federal Title IX

Students are entitled to safe, equitable circumstances, free of harassment, discrimination, or interpersonal violence. If you find yourself in a violating situation, you are encouraged to seek support through the Rice SAFE Office. Under Federal regulations, most Rice employees, including me, are required to report any incident of non-consensual interpersonal behaviors that is disclosed to them. That report is to be made to campus Title IX professionals, who can intervene in support of the student. For more information, visit https://safe.rice.edu or email titleixsupport@rice.edu.

Use of Machines in the Classroom

Last, but not least …

Cell phones must be turned off during class unless I ask you to use them for class purposes. If you anticipate the need to take an urgent call, please set your phone to vibration and take the call outside the room. Laptops or other small devices may be used only for specific class purposes such as participating via Zoom or taking notes. I strongly recommend limiting the use of computers during lecture as it can be distracting to you as well as to your neighbors!

In case of remote instruction (should it become necessary), please have live camera presence in class if at all possible, or at least a static picture of you in Zoom, and be available for discussion. I will expect students to be punctual, to dress like in the regular classroom, and to refrain from eating or other distractions during class.