CMSC 727 Neural Computation Spring 2022

Time and Place: Tuesdays and Thursdays, 12:30 pm – 1:45 pm, ONLINE

Instructor: James Reggia myLastName AT umd DOT edu
Office Hours: Thursdays, 2 pm – 3 pm, or by appointment

Teaching Asst.: Hanyu Wang hywang66 AT umd DOT edu
Office Hours: Mon. 11 am – noon, Weds. 2 pm – 3 pm, or by appointment

Prerequisites: Graduate standing or permission of instructor. We assume that you have a basic background in undergraduate calculus, linear algebra, elementary probability theory, and Python programming. No background in neurobiology/neuroscience is assumed.

Class web page: https://www.cs.umd.edu/class/spring2022/cmsc727/
Gives exam dates, homework assignments and their due dates, lecture slides, reading assignments, links to online class sessions and office hours, and links to other useful information.

Content:

The field of neural computation involves the study of representations and algorithms that are inspired by information processing in the brain. This field has been growing very rapidly, reflecting its successful applications in machine learning and increasing interest in its potential to produce an artificial mind. In this course we will systematically cover the most important neural computation methods, illustrate their use for machine learning and for cognitive/brain modeling, provide hands-on experience training neural models, and introduce active research and application areas such as deep learning and programmable neural networks. The enormous amount of information currently available about neural networks means that we will be selective in what we cover this semester. In doing this we will emphasize the role of neural computation in machine learning and AI, although some multi-disciplinary information (cognitive science, neuroscience, etc.) will also be presented. By the end of the semester you should be familiar with many of the central concepts and commonly used neural network methods as outlined below.

1. *Basic concepts*: how the brain computes; neural network architectures, activity dynamics and learning rules; self-organization and emergent intelligence; competition and cooperation in networks, local vs. distributed representations, historical perspective, applications.

2. *Basic feedforward architectures using supervised learning*: logical neurons, perceptrons, delta rule learning with linear associative memory, gradient descent, logistic regression, and error backpropagation. Some core concepts from machine learning (data standardization, performance assessment, overfitting, regularization, etc.) will be reviewed along the way.

3. *Basic recurrent architectures using Hebbian learning*: Hebb's rule, linear autoassociative memories, neural nets as dynamical systems, attractor networks (e.g., Hopfield nets, BSB), energy functions, simulated annealing, Boltzmann machines, restricted Boltzmann machines.

4. *Deep learning*: brain inspiration, Neocognitron, deep convolution networks, generative adversarial networks, deep belief networks, deep Boltzmann machines, deep autoencoders, machine vision & natural language processing applications.

5. *Other feedforward methods*: error backpropagation variations/extensions (e.g., RPROP), interpreting trained networks, radial basis function networks.

6. *Contemporary architectures using supervised learning*: echo state networks, recurrent backpropagation networks, backpropagation through time, long short-term memory (LSTM), attention mechanisms, transformer networks, hyperdimensional computing.

7. *Unsupervised or no learning*: competitive learning, self-organizing feature maps, adaptive resonance theory, continuous energy-minimizing attractor nets, oscillatory neural nets.

8. *Reverse engineering the brain*: neurobiological basis of cognition, large-scale brain models, brain-inspired neurocognitive architectures, programmable neural networks, neural virtual machines, NeuroLisp, strong AI and the artificial consciousness movement, prospects for a technological singularity.

9. *Related topics as time permits*: ensemble methods, neural nets in reinforcement learning, swarm intelligence and its relation to neural networks, evolutionary computation and neural networks, learning of activation rules, applications, software packages, hardware implementations and neuromorphic systems, large-scale networks, spiking neurons.

Workload and Grading: There will be regular reading and homework assignments. Assignments will typically include some online work using Python and related code. Assignments and examinations are always to be treated as independent work. A more open-ended and substantial small group "semester project" is required. Grading will be based on homework assignments, quizzes, and class participation (collectively 20%), semester project and presentation (30%), midterm (25%), and second exam (25%).

Reading Sources:

Charu Aggarwal, *Neural Networks and Deep Learning*, Springer, 2018. ISBN: 978-3-319-94462-3 will serve as the primary textbook with regular assigned readings.

R. O'Reilly, et al., *Computational Cognitive Neuroscience*, 4th Edition, 2020.
A free pdf version of this book is available at: https://github.com/CompCogNeuro/ed4

I. Goodfellow, Y. Bengio, A. Courville, *Deep Learning*, MIT Press, 2016. A free pdf version of this book is available online at http://www.deeplearningbook.org

Additional readings from the literature will be posted as pdf files on the class web site.

Examinations: Exam 1: Thursday, March 17
Exam 2: Tuesday, May 10
Final Exam Slot: Tuesday, May 17, 1:30 pm – 3:30 pm (reserve this time)

Disabilities: Any student eligible for and requesting reasonable academic accommodations due to a disability needs to provide the instructor with a letter of accommodation from the Office of Disability Support Services within the first two weeks of the semester.

Class Absence Policy: The campus has a policy governing class absences. This policy requires instructors to provide the following information. For CMSC 727, the "major scheduled grading events" are the two exams and the semester project. A maximum of one self-signed medical excuse for late submission of other grading events will be accepted.

Academic Integrity: All homework assignments and examinations are to be done individually and independently unless specifically stated otherwise in writing; all submitted work must be your own. All students are expected to be familiar with and to uphold the Code of Academic Integrity (see https://studentconduct.umd.edu/you/students). Posting homework solutions online violates academic integrity policy. Further details of CS Dept. Academic Integrity policies are at http://www.cs.umd.edu//class/resources/academicIntegrity.html