Course Syllabus

Instructor: Dr. Sudeep Pasricha (sudeep@colostate.edu), ENGR B119

Lectures: Tu/Thu 12:30pm – 1:45pm, BHSCI 107 (online recordings available via Echo360)

Office Hours: 9am – 10:30am, Fridays

Course TA: Joydeep Dey (joydeep.dey@colostate.edu)

TA Hours: TBD

Course Description: Machine learning is becoming pervasive in embedded computing platforms, such as smart mobile systems, wearable IoT devices, and autonomous vehicles. This course will present recent advances towards the goal of enabling efficient implementation of deep machine learning models on embedded systems. Specifically, it will provide an overview of 1) the theoretical foundations and motivations behind various deep learning models, 2) software modeling and optimization techniques for these models, 3) hardware platforms and architectures to support efficient execution of machine learning models, and 4) hardware-software co-design approaches for machine learning. The course will cover emerging machine learning models, custom hardware accelerators, as well as paradigms including processing-in-memory, memristors, and photonics for machine learning. The course is very topical and relevant for graduate and senior undergraduate students in computer engineering, computer science, data science, and electrical engineering, as well as practitioners in industry and students in other engineering departments who are interested in data engineering, machine learning, and embedded systems. The course does not assume prior expertise in machine learning; however knowledge of Python is essential. Some background in computer architecture (ECE452 or similar) would also be helpful. Students will get hands-on experience in the design and optimization of deep machine learning models, as well as the analysis of these models on emerging hardware platforms. This is a unique course that is well suited for beginners and experts alike, who want to comprehend the state-of-the-art in deep machine learning software and hardware design, and understand future trends and opportunities in this exciting field.

Prerequisites: ECE251 or CS270 or equivalent computer organization course; programming experience in Python

Textbook: None. The course will cover materials from various books and conference/journal articles.

Syllabus: Here is a tentative outline and syllabus for this course.

Module	Weeks	Lecture Content
1	1-2	Intro to Machine Learning (ML) and Embedded Computing
2	3-4	Deep Neural Networks (DNNs/CNNs)
3	5-6	Software Optimizations for ML
4	7-8	Hardware Acceleration for ML
5	8	HW/SW Co-design
6	9	RNNs, LSTMs, and GRUs
7	10	Unsupervised Learning
8	11	Anomaly Detection and Security
9	12	Autoencoders and GANs
	13	Fall Recess
10	14	Advanced ML Models
11	15	Emerging Directions and Applications in ML
12	16	Final Week: Project Deliverables

Grading: The final grade will be on a curve. Grading is based on the following components:

- Homework/Lab Assignments (5): 30%
- Reading Assignments (10): 20%
- Quizzes (2): 10%
- Class Participation: 10%
- Final Project: 30%
  - Final presentation: 10%
  - Project report: 20%

Grading Scale:

>95%	90-94%	85-89%	80-84%	75-79%	70-74%	65-69%	55-64%	40-55%	<40%
A+	A	A-	B+	B	B-	C+	C	D	F

Assignments: Homework assignments will involve working with Python and Tensorflow/Keras, as well as tools for embedded platform exploration. Reading assignments will involve reading technical research papers and summarizing their key contributions and a critique in around 500 words.

Submission Policy: Homework and reading assignments will be assigned throughout the semester. You are allowed late submission up to 3 days on one homework and one reading assignment. You can also skip one reading assignment of your choice, without impacting your grade. Otherwise all homework and reading assignments should be submitted before the deadline via Canvas, and ***late submissions will not be graded!***

Re-grading Policy: Re-grading requests should be made within a week from the date of the graded item (homework, exam, or project) becoming available.

Academic Integrity: All submitted work should be your own. Copying of language, structure, images, ideas, or thoughts of another, and representing them as one's own without proper acknowledgement (from github code repos, other web sites, books, papers, other students, etc) and failure to cite sources properly is not acceptable. Sources must always be appropriately referenced, whether the source is printed, electronic, or spoken. Minor first infraction in HWs and presentations will lead to a zero score + one letter level (e.g. A to B) reduction in course grade. Project or Major or repeated infractions in HWs and presentations will result in “F” grade for the course + report to Dean’s Office. For more information see CSU’s Academic Integrity Policy: https://tilt.colostate.edu/AcademicHI (Links to an external site.) and Student Conduct Code: https://resolutioncenter.colostate.edu/student-conduct-code/ (Links to an external site.)

Attendance: I encourage everyone to attend all the lectures and actively participate in class discussions.

Appointment: I encourage you to make at least one appointment with me during the semester for advice or to discuss research opportunities, independent study, research ideas, course suggestions, concerns, or any other topic you feel is appropriate.