ESE3060 Deep Learning: A Hands-on Introduction

Course Number & Title (A.1)	**ESE3060** Deep Learning: A Hands-on Introduction
Credit Units (A.2)	1 CU (3 semester hours)
Class/Laboratory Schedule	Lecture: 3 hrs/week
Instructor (A.3)	Hamed Hassani <hassani@seas.upenn.edu>
Text(s)/Required Materials (A.4)	We will use several resources that are available online, including the following textbook: Dive Into Deep Learning, Zhang, Lipton, Li, Smola (https://d2l.ai/)
Catalog Description (A.5a)	This course will serve as an introductory and hands-on dive into the area of deep learning. The main goal is to to educate the students on (i) the commonly-used neural network architectures and proficiency in training them, (ii) Some of the main problems that deep learning systems have successfully addressed (formulation, architecture, data sets, etc). There will be no theory in this course. After finishing this course, the students should be very comfortable with pytorch programming as well as training deep learning models.
Prerequisites (A.5b)	The students should be comfortable with python programming in order to register for the course. A basic understanding of supervised learning will be helpful. ESE 2000, 2240 or any basic AI-related course
Course Satisfies (A.5c)	[ ] Math [ ] Science [**X**] Engineering [ ] Technical Elective [ ] TBS **(check only one, UG curric impact only)** **CMPE** [ ] Required [ ] Selected Elective [ ] Elective **(check only one)** **EE** [ ] Required [ ] Selected Elective [**X**] Elective **(check only one)** **SSE** [ ] Required [ ] Selected Elective [**X**] Elective **(check only one)** **Engineering Design Component:** 50% Note: ABET EAC Criterion 5b: Engineering design is the process of devising a system, component, or process to meet desired needs. It is a decision-making process (often iterative), in which the basic sciences, mathematics, and the engineering sciences are applied to convert resources optimally to meet these stated needs.
Course Web	Canvas site will be established.
Course Outcomes (A.6a)	1 understand various scenarios of machine learning and how deep learning can help to address the corresponding challenges. 2a ability to train deep neural architectures using pytorch programming. 2a ability to identify the appropriate neural architecture depending on the problem and data 2c ability to incorporate social and economic considerations when training large-scale models 3 demonstrate ability to communicate deep learning design and function to students from a broad range of science and engineering disciplines 4 ability to understand the ethical and issues of deep learning design and deployment in various application domains 5 ability to work effectively in a team to apply deep neural networks to solve problems 6 ability to develop and conduct appropriate optimization techniques and experimentation, analyze and interpret data, and use engineering judgment to draw conclusions. 7 understand and use open-source code to carry out hands-on labs and team-based project
Contribution towards Program Outcomes (A.6b)	1 – high 2 – high 3 – medium 4 – low 5 – medium 6 – high 7 – medium
Topics Covered (A.7)	

Weekly/Session Schedule (A.7)	Week 1: Introduction to pytorch Week 2: Basics of machine learning, loss functions, data sets, optimization for deep learning Week 3: Automatic differentiation, stacking layers, MLPs Week 4: Automatic differentiation, stacking layers, MLPs Week 5: Convolutional neural networks Week 6: Convolutional neural networks (feature extraction, network visualization) Week 7: Training Neural Networks: Data Augmentation, Hyperparameter Tuning, Initialization, Choice of activation function and architecture Week 7: Dropout, batchnorm, Resnets Week 8: Autoencoders and generative models (GANs, VAEs, Diffusion models) Week 9: Autoencoders and generative models (GANs, VAEs, Diffusion models) Week 10: Recurrent neural networks Week 11: Attention and transformers Week 12: LLMs, prompting, in-context learning Week 13: Projects and presentations Week 14: Projects and presentations
Ethical Issues and Treatment	The course will briefly touch on ethical issues of deploying machine learning models in practice, including security and privacy concerns
Grading Details	70% Homework/Lab; 30% Team Project/Presentation
Prepared By/Date	Hassani, March, 2024

Deep Learning: A Hands-on Introduction

Goals

- Learn to create sophisticated deep learning models
- Understand different deep learning architectures; when/why/how to use them
- Learn by doing (minimal theory, no exam, HWs/project, completely hands-on)

Syllabus

Tensors/Pytorch

Convolutional Neural Networks

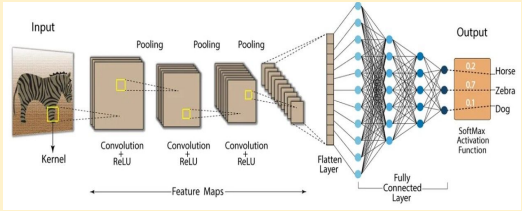

Autoencoders

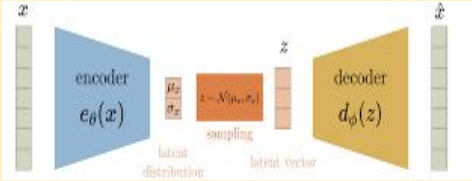

Diffusion Models

DALL·E

Recurrent Neural Networks

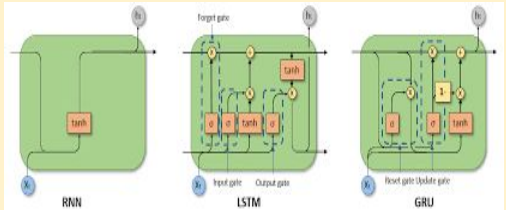

Transformers

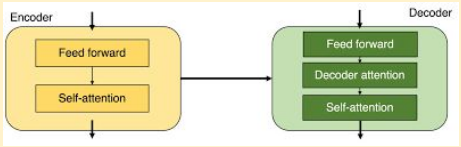

ChatGPT

Prerequisites

Very basic knowledge in ML/AI

Python, or any other PL

Desire to learn and create!