MULTIMEDIA INFORMATICS

MMI 727– Deep Learning: Methods and Applications

SYLLABUS

Year, Semester:	2023-2024 Fall
Course Conduct:	Face-to-face Lectures II-05 Tuesday@13:40 Lecture videos on YouTube & Lecture notes on ODTUClass Lab sessions on Google Colab Students are expected to watch the lecture videos and study the course material before attending the weekly synchronous sessions. There will be quizzes from the content and attendance is expected.
Lecturer:	Prof. Dr. Alptekin Temizel, atemizel@metu.edu.tr
Teaching Assistant:	Ayberk Aydın, aayberk@metu.edu.tr

Course Objective

This course aims to give background knowledge on several topics related to deep learning and provide a laboratory environment for practical applications. Backpropagation, convolutional neural networks, recurrent networks, Long Short-term memory (LSTM), generative adversarial networks, attention and transformers are some of the core topics that will be covered through the lectures. The course aims to balance theory and practice in that it will involve students implementing various algorithms, testing those algorithms under several domains and accessing GPU clouds during laboratory sessions to code the program examples using PyTorch.

Textbook:

I. Goodfellow, Y. Bengio, and A. Courville. Deep Learning. MIT Press, 2016.

Reference Material:

CS231n: Convolutional Neural Networks for Visual Recognition - http://cs231n.stanford.edu/
Deep Learning Resources - https://deeplearning4j.org/deeplearningpapers

Reading Material:

How Smart Machines Think. MIT Press. Sean Gerrish, Kevin Scott.
The Cambridge Handbook of Artificial Intelligence, Keith Frankish, William M. Ramsey.

Attendance and Absences

Grade Distribution:

Lab reports and 3x Assignments	50%
Final project	25%
5x Quizzes	25%

Deliverables

Documents and necessary files of the assignments must be uploaded to ODTUClass by students before the specified due dates.

Late Submissions

Late submissions are accepted with a penalty of 10% grade decrease for each day delayed.

University Policies

All students are **expected to obey** the university code of integrity and avoid academic dishonesty or plagiarism.

No	Date	
1	3 October	**Introduction to the Course and Fundamentals of Deep Learning**
2	10 October	**Introduction to Machine Learning** Introduction to Machine Learning Introduction to Neural Networks Deep Learning Basics and Architecture
3	17 October	**Introduction to Deep Learning** History of Deep Learning and Convolutional Networks Deep Learning Ecosystem Linear Machines Feature Learning, Trainable Features, Hierarchical Feature Representation, Invariant Features
4	24 October	**Lab Session – I** Deep Learning environment, setups and hardware Cloud systems and Colab environment Creating a neural network from scratch Hands-on MNIST handwritten digit recognition example (PyTorch/Keras)
5	31 October	**Convolutional Neural Networks-I** Convolutional Network Model Convolution and Pooling Operations, Activation Functions
6	7 November	**Convolutional Neural Networks-II** Object Recognition/Image Classification AlexNet, Data Augmentation, Dropout Inception Module, Resnets Network Design Guidelines
7	14 November	**Convolutional Neural Networks-III** Object Detection, R-CNN, Fast/Faster R-CNN, YOLO Internal Covariate Shift, Batch Normalization Initialization Issues, Xavier Initialization
8	21 November	**Lab Session -II**
9	28 November	**Learning with Memory - I** Time-delay Neural Network (TDNN) Recurrent Neural Networks (RNN) Character-Level Language Model Image Captioning, Attention
10	5 December	**Learning with Memory – II** Natural Language Processing, Word2Vec Shortcomings of RNNs, Long-Short Term Memory (LSTM) Self-Attention and Transformers
11	12 December	**Generative Models** PixelRNN/CNN Variational Autoencoders Generative Adversarial Networks Diffusion Models
12	19 December	**Optimization for Deep Learning** Gradient Descent Momentum and Nesterov accelerated gradient Gradient descent variants, Adagrad, Adadelta, RMSProp, Adam
13	26 December	**Applications** Image Segmentation Low-cost/mobile Networks
14	2 January	**Deployment and Network Optimization** Model Pruning and Quantization Aware Training Deployment Platforms Tools for Deployment and Network Optimization

There’ll also be extra content during face-to-face classes to cover recent developments in the field and go through some case studies such as:

- Deep Neural Nets: 33 years ago and 33 years from now (Andrej Karpathy), Implementation and Comparison of original LeNet mode
- Closer look at ResNet and Inception architectures
- Recent architecture on augmentation and regularization
- Recent CNN and transformers based object detection methods
- Implementation of Attention and Transformer based models
- Hyperparameter & Model Architecture Tuning
- Recent segmentation models