	U-LAS30 20036 LE10			
< >	Fundamentals of Artificial Intelligence-E2 Fundamentals of Artificial Intelligence-E2			LALA, Divesh Kanu
		( )	( )	
		2	1	
	2025	5		

[   ]

Society is currently in the middle of an "AI boom", with recent computers seemingly able to perform tasks which would have been very difficult even a few years ago. In this course we explain what enables these machines to successfully learn like a human does. We start from discussing fundamental machine learning models such as linear regression to more recent breakthrough architectures such as transformer models. Importance will also be placed on how to properly train and evaluate AI models. We then apply what we have learned to understand how the models are used to make a conversational robot.

[   ]

Students will gain an understanding of machine learning techniques and architectures and also how to construct basic models. They will also know how to properly evaluate AI models using real data.

[   ]

1. Artificial intelligence (1 week)
We will discuss the current AI boom. What does it mean for society? Can we say that AI is truly "intelligent"? We will then explain how machine learning can be linked to concepts of human intelligence.

2. Data analysis (2 weeks)
Collecting data is the first step in training machine learning models, so we will discuss how a wide range of data can be collected, either by using publicly available corpora or collecting it ourselves. We also show how to get quality data to match the goal of our AI models.

3. Machine learning (4 weeks)
In this set of lectures we focus on the fundamental concepts of machine learning. We review basic mathematics required to study machine learning, including vectors and matrices. We then begin with simple supervised linear and logistic regression models to introduce concepts in machine learning and apply these models to a real data set to understand how the models work in practice. Unsupervised models such as clustering are also discussed. Emphasis will be placed on the correct way to train models.
These lectures will be accompanied by basic programming in Python

4. Deep learning (2 weeks)
Once the concepts of machine learning are understood, we move on to more state of the art methods involving deep learning and transformer models, which are responsible for the breakthroughs in the current AI boom. We describe the basic architecture of these models and how they can be used to create powerful classifiers for a range of different data types.

5. Evaluation methods (1 week)
We take a special look at the evaluation of AI models. Common mistakes related to bias and evaluation

Fundamentals of Artificial Intelligence-E2(2)

metrics will be discussed to reinforce good practice when testing and evaluating models.

6. Application (4 weeks)
Armed with the knowledge of the previous lectures, we discuss how AI models can be used in the context of human conversation. We focus on large language models such as ChatGPT and human conversation behavioral models which can be used to produce spoken dialogue systems. Various language, audio and visual models will also be described. As part of these lectures students will choose a simple classification task related to human conversation and then successfully construct an AI model.

7. Feedback (1 week)

[   ]

[   ]

Attendance and participation (20%), exercises (60%) and a final report (20%).

[   ]

Handouts

[   ]

Students should aim to review course content for 30 minutes before and after class and if possible practice programming on their own machines using Python.

[   ]