

INDEX OF AI TOPICS & TERMS IN THIS COURSE



LUKA ANICIN
CHRIS HAROUN

Please Note: Topics are sorted alphabetically. Below each topic is the definition and where to find the topic in the course and in the book version of the course.

AGI or Artificial General Intelligence

Definition:

Artificial General Intelligence (AGI) is also called Strong AI. The goal of AGI is to mimic human intelligence. This type of AI can learn, reason, and adapt to new situations on its own. There really isn't an example of AGI that exists yet and there are plenty of science fiction movies that have examples of strong AI from the movie 2001 a Space Odyssey to one of my favorite video games called Detroit become human. AGI is the scientific equivalent of the splitting of the atom.

Section Topic is Covered in the Course and Book Version of the Course:

Section 2: Intro. to Artificial Intelligence (Machine + Deep Learning & More)

Lecture or Chapter Name Topic is Covered in the Course and Book Version of the Course:

*[A/B/C] 3 Categories of AI: Narrow (ANI), General (AGI) & Super (ASI)

ANI or Artificial Narrow Intelligence

Definition:

Artificial Narrow Intelligence (ANI) is designed for relatively simple tasks from an AI perspective like basic voice based

technologies, self-driving cars, and movie recommendation engines.

Section Topic is Covered in the Course and Book Version of the Course:

Section 2: Intro. to Artificial Intelligence (Machine + Deep Learning & More)

Lecture or Chapter Name Topic is Covered in the Course and Book Version of the Course:

* [A/B/C] 3 Categories of AI: Narrow (ANI), General (AGI) & Super (ASI)

ASI or Artificial Super Intelligence

Definition:

Artificial Super Intelligence (ASI), which means that AI would be smarter than humans in almost every way. While this concept is speculative, it's widely discussed among AI researchers and futurists.

Section Topic is Covered in the Course and Book Version of the Course:

Section 2: Intro. to Artificial Intelligence (Machine + Deep Learning & More)

Lecture or Chapter Name Topic is Covered in the Course and Book Version of the Course:

* [A/B/C] 3 Categories of AI: Narrow (ANI), General (AGI) & Super (ASI)

AutoML

Definition:

AutoML is a process that automates the development and deployment of Machine Learning models. It involves using AI algorithms and techniques to automatically select the best model architecture which is called Automatic Model Selection. This process automatically evaluates and selects the most suitable Machine Learning algorithm for a given problem or dataset, based on performance metrics and a goal or goals that you set.

Section Topic is Covered in the Course and Book Version of the Course:

Section 2: Intro. to Artificial Intelligence (Machine + Deep Learning & More)

Lecture or Chapter Name Topic is Covered in the Course and Book Version of the Course:

* [A] AutoML and Transfer Learning

Computer Vision

Definition:

Computer Vision is a field of Artificial Intelligence that enables computers and machines to interpret and understand visual information from the world, such as images, videos, or real-time camera feeds. It's like teaching a computer to "see" and make sense of the visual world, just as humans do.

Section Topic is Covered in the Course and Book Version of the Course:

Section 2: Intro. to Artificial Intelligence (Machine + Deep Learning & More)

Lecture or Chapter Name Topic is Covered in the Course and Book Version of the Course:

* [A/B] Computer Vision (Using Tesla as an Example)

DL or Deep Learning

Definition:

Deep Learning (DL) is a subset of Machine Learning (ML). Deep Learning uses Artificial Neural Networks to model and solve complex problems. It's inspired by the structure and function of the human brain. Deep Learning is more complex than Machine Learning and requires better hardware.

Section Topic is Covered in the Course and Book Version of the Course:

Section 2: Intro. to Artificial Intelligence (Machine + Deep Learning & More)

Lecture or Chapter Name Topic is Covered in the Course and Book Version of the Course:

* [A/B] A.I. vs Machine Learning vs. Deep Learning

LLM or Large Language Model

Definition:

A Large Language Model (LLM) like ChatGPT is trained using a massive amount of text in order to provide value to the user (i.e., answering a ChatGPT questions or prompt).

Section Topic is Covered in the Course and Book Version of the Course:

Section 3: ChatGPT

Lecture or Chapter Name Topic is Covered in the Course and Book Version of the Course:

* [A/B/C] Your First Prompt Exercise - Directional Prompting

ML or Machine Learning

Definition:

Machine Learning (ML) is a subset of AI that involves developing algorithms to learn patterns from data. Machines learn from experience, just like humans do.

Section Topic is Covered in the Course and Book Version of the Course:

Section 2: Intro. to Artificial Intelligence (Machine + Deep Learning & More)

Lecture or Chapter Name Topic is Covered in the Course and Book Version of the Course:

* [A/B] A.I. vs Machine Learning vs. Deep Learning

NER or Named Entity Recognition

Definition:

Named Entity Recognition or NER is a Natural Language Processing (NLP) task and it involves identifying and

classifying specific entities, such as names of people, organizations, locations, or dates, within a given text.

Section Topic is Covered in the Course and Book Version of the Course: Section 2:

Intro. to Artificial Intelligence (Machine + Deep Learning & More)

Lecture or Chapter Name Topic is Covered in the Course and Book Version of the Course:

* [A] Natural Language Processing (NLP) and NLP Applications

NLP or Natural Language Processing

Definition:

Natural Language Processing, also called NLP, is a branch of Artificial Intelligence that focuses on enabling computers to understand, interpret, and generate human language in a way that is both meaningful and useful.

Section Topic is Covered in the Course and Book Version of the Course:

Section 2: Intro. to Artificial Intelligence (Machine + Deep Learning & More)

Lecture or Chapter Name Topic is Covered in the Course and Book Version of the Course:

* [A] Natural Language Processing (NLP) and NLP Applications

Prompt Engineering

Definition:

Prompt Engineering is an Artificial Intelligence term used when discussing Natural Language Processing (i.e., asking an AI product a question).

Section Topic is Covered in the Course and Book Version of the Course:

Section 3: ChatGPT

Lecture or Chapter Name Topic is Covered in the Course and Book Version of the Course:

* [A/B/C] Introduction to ChatGPT (What, Why and How of Section 3)

Reinforcement Learning

Definition:

Reinforcement Learning is a subset of Machine Learning, where an agent like a computer or a robot learns to make decisions by interacting with an environment. The agent receives feedback in the form of rewards or penalties, and its goal is to maximize the cumulative reward over time.

Section Topic is Covered in the Course and Book Version of the Course:

Section 2: Intro. to Artificial Intelligence (Machine + Deep Learning & More)

Lecture or Chapter Name Topic is Covered in the Course and Book Version of the Course:

* [A] Supervised & Unsupervised Machine Learning + Reinforcement Learning

Role Prompting

Definition:

When prompting an LLM (Large Language Model) like ChatGPT, at the beginning of the prompt, we can add a part we call `System` prompt or a Role prompt. A Role prompt tells the LLM who to be when answering a question. For example, "you are a professional soccer player, Act as a senior lawyer partner, etc." By defining a role, we've given the AI more context and set expectations for the quality and tone of the response.

Section Topic is Covered in the Course and Book Version of the Course:

Section 3: ChatGPT

Lecture or Chapter Name Topic is Covered in the Course and Book Version of the Course:

* [A/B/C] Role Playing

Supervised Learning

Definition:

Supervised Learning is a subset of Machine Learning, where labeled data sets (like columns in a spreadsheet) are provided as inputs and the output is derived (i.e., through a Regression Analysis - meaning here is X, solve for Y).

Section Topic is Covered in the Course and Book Version of the Course:

Section 2: Intro. to Artificial Intelligence (Machine + Deep Learning & More)

Lecture or Chapter Name Topic is Covered in the Course and Book Version of the Course:

* [A] Supervised & Unsupervised Machine Learning + Reinforcement Learning

System Prompting

Definition:

When prompting an LLM (Large Language Model) like ChatGPT, at the beginning of the prompt, we can add a part we call `System` prompt or a Role prompt. A Role prompt tells the LLM who to be when answering a question. For example, "you are a professional soccer player, Act as a senior lawyer partner, etc." By defining a role, we've given the AI more context and set expectations for the quality and tone of the response.

Section Topic is Covered in the Course and Book Version of the Course:

Section 3: ChatGPT

Lecture or Chapter Name Topic is Covered in the Course and Book Version of the Course:

* [A/B/C] Role Playing

Tensor Processing Unit

Definition:

TPU stands for Tensor Processing Unit, which was invented by Google for use in AI and in other applications. The great thing about TPUs is that they are designed in a way to use less processing power than GPUs or CPUs. A TPU can be used with Google Cloud services, which is Google's version of Amazon Web Services, meaning cloud computing.

Section Topic is Covered in the Course and Book Version of the Course:

Section 2: Intro. to Artificial Intelligence (Machine + Deep Learning & More)

Lecture or Chapter Name Topic is Covered in the Course and Book Version of the Course:

* [A/B] AI Hardware and Infrastructure: Enabling Scalable and Affordable AI

TensorFlow**Definition:**

A free and open source library to help us develop AI algorithms, which was developed by Google.

Section Topic is Covered in the Course and Book Version of the Course:

Section 2: Intro. to Artificial Intelligence (Machine + Deep Learning & More)

Lecture or Chapter Name Topic is Covered in the Course and Book Version of the Course:

* [A/B] AI Hardware and Infrastructure: Enabling Scalable and Affordable AI

TPU

Definition:

TPU stands for Tensor Processing Unit, which was invented by Google for use in AI and in other applications. The great thing about TPUs is that they are designed in a way to use less processing power than GPUs or CPUs. A TPU can be used with Google Cloud services, which is Google's version of Amazon Web Services, meaning cloud computing.

Section Topic is Covered in the Course and Book Version of the Course:

Section 2: Intro. to Artificial Intelligence (Machine + Deep Learning & More)

Lecture or Chapter Name Topic is Covered in the Course and Book Version of the Course:

* [A/B] AI Hardware and Infrastructure: Enabling Scalable and Affordable AI

Transfer Learning

Definition:

Transfer Learning is when big companies such as Google, Microsoft, Amazon, IBM and OpenAI train huge models for months in their datacenters and then open source the models for public use. These models are trained on general purpose data, with a goal to learn basic constructs like language, shapes, sizes etc.

Then, we can take these models and adjust them to work on our specific problems in our companies. And instead of training those models for months, we need just a couple of hours to train it on our custom data!

Section Topic is Covered in the Course and Book Version of the Course:

Section 2: Intro. to Artificial Intelligence (Machine + Deep Learning & More)

Lecture or Chapter Name Topic is Covered in the Course and Book Version of the Course:

* [A] AutoML and Transfer Learning"

Unsupervised Learning

Definition:

Unsupervised Learning is a subset of Machine Learning where a computer is given a lot of data and it is asked to find patterns or clusters. For example, these types of customers prefer these types of products.

Section Topic is Covered in the Course and Book Version of the Course:

Section 2: Intro. to Artificial Intelligence (Machine + Deep Learning & More)

Lecture or Chapter Name Topic is Covered in the Course and Book Version of the Course:

* [A] Supervised & Unsupervised Machine Learning + Reinforcement Learning