Lecture 1
Some Machine Learning Preliminaries

By
Dr. Muhammad Alrabeiah
Electrical Engineering Dept., KSU
Spring 2022

January 12th, 2022

Disclaimer These notes are still under development, and they have not been subjected to proper review and revision.

Contents

Chapter 1

Basic Definitions

Suggested readings:
Chapters 1 and 5 of "Deep Learning," by Ian Goodfellow et al, 2016.

Machine learning and artificial Intelligence (AI) are closely interwind subjects, to the point that they are very often confused to be synonyms. Therefore, the first step in our journey through this course will be establishing some intuitive and working understanding of the difference between the two. Emphasis here is on the adjectives "intuitive" and "working," for a formal discussion on the difference between the two could take us down a philosophical and historical rout, which is out of the scope of this undergraduate-level course. Once the difference is clear, we will move towards laying some necessary groundwork for neural networks and deep learning. We will develop definitions for some basic principles of machine learning, supervised learning, task, experience, and performance measure to name four. Then, we will discuss some fundamental tasks in supervised machine learning, in particular regression and classification.

1 What is AI? And How Is It Different From Machine Learning?

Complex systems such as a robot or a self-driving vehicle (e.g., a terminator and Audi RSQ if you are a Sci-Fi movie geek) are good examples of AI. They are capable of acquiring data from their environments using various types of sensors (e.g., RGB cameras, radars, LiDARS,...etc),

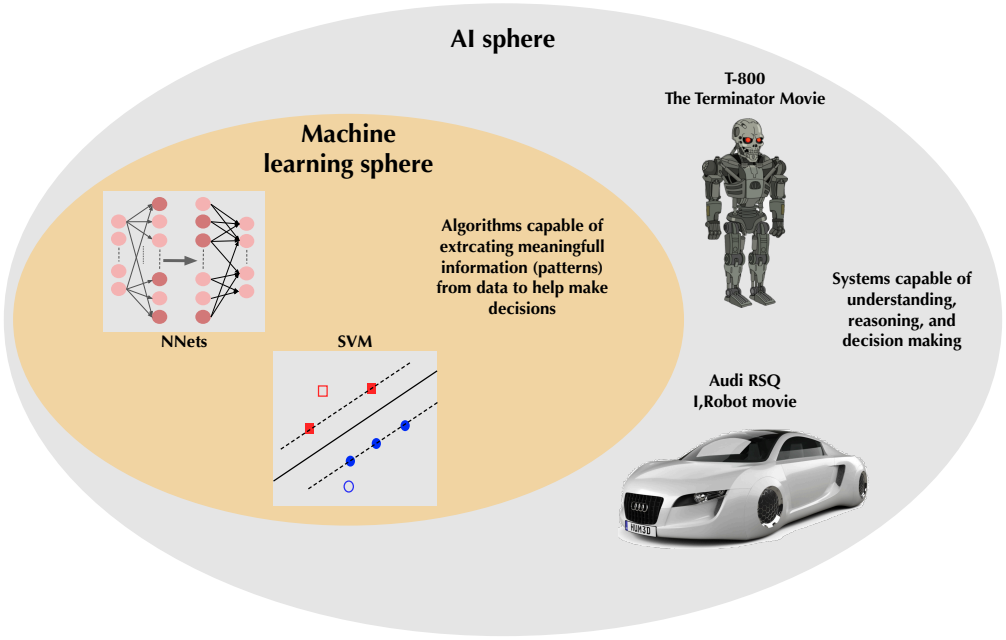

Figure 1.1: Visualization of AI and machine learning and their relation.

analyzing those data to understand their environments, and making decisions to interact with their surroundings. All those capabilities define what could be considered the three pillars of AI, which are: (i) sensing, (ii) understanding, and (iii) decision making. Hence, throughout this course, we adopt a definition for AI that is based on those three pillars.

Definition 1.1. Artificial Intelligence is the field concerned with developing intelligent systems. Such systems are commonly composed of hardware and software components that enable the acquisition and analysis of sensing data for the purpose of: (i) understanding the environment where those systems function; and (ii) making decisions that help those systems interact with their surroundings.

It is important to emphasize here that the definition above is not meant to be comprehensive and rigorous, but simple and intuitive. It should be good enough to help us differentiate AI from machine learning. More comprehensive discussion and rigorous definition could be found in AI textbooks like [1].

For an AI system to understand its environment and make decisions, it needs to extract useful information from the data it acquires. This specific need is what machine learning aims to satisfy. It provides ways to analyze acquired data and "recognize" patterns relevant

to what the AI wants to do. This may seems a bit vague, so we will give a clear definition for machine learning that highlights its important components.

Definition 1.2. Machine learning is the field concerned with developing "learning algorithms" that recognize certain patterns from the data they are given to perform a well-defined task.

With the above definition, we can now develop a good understanding to what the relation between AI and machine learning is; *AI is the sphere where machine learning exists, and machine learning is an approach by which AI can function,* see Figure 1.1. Note that AI can function without machine learning; however, the last decade of AI research has shown that machine learning is the best known approach to help AI realize its objectives.

2 What Is A Learning Algorithm?

A keywords stand out in Definition 1.2 for the keen reader, which is learning algorithm. We are going to provide a definitions for that keyword along with some other relevant concepts. From textbook [2], we will adopted the following definition for a learning algorithm

Definition 2.1. Learning algorithm: A computer algorithm is said to learn from experience E with respect to some class of tasks T and performance measure P, if its performance in task T as measured by P improves with experience E.

This definition introduces several important machine learning concepts. Below is a breakdown of those concepts with informal definitions.

- **Task:** This word defines what is expected from the learning algorithm to do. To understand this, let's take the robot in Figure 1.1 as an example. This robot might need to visually understand its surrounding. Hence, it may deploy a camera to acquire color images about its environment (like how we human observe our surrounding). Now, one way a learning algorithm might help that robot is through identifying objects around it, e.g., dog, cat, couch, table, cup, door,... etc. Those objects, despite how diverse they could be, define a discrete yet large set, which we will call the set of "object classes." The learning algorithm is expected to see the input image and assign each object in that image an object class, and that process is exactly what we mean by a task. In the

realm of machine learning, different tasks often belong to a common category of tasks. In our example here, the task of assigning an object a class falls under the "classification" category of tasks. More on that later.

- **Performance measure:** Because a learning algorithm is expected to perform a specific and well-defined task, its performance in doing that task needs to be assessed with a form of a measure. For the robot example above, the "accuracy" of its object-to-class assignment needs to be measured with a function we will call performance metric. This metric allows fair comparison between learning algorithms that attempts to perform the same task. One good example that we will encounter soon is the "cross-entropy" function. This is a non-negative function that measures how accurate the classification of the algorithm is. We are getting ahead of ourselves here, so let's leave the discussion on cross-entropy to another lecture.

- **Learning Experience:** This is a very important concept in machine learning, if not the most important. The experience encompasses how the algorithm learns and what it learns from. In simple words, it is the set of data points or examples—henceforth called *dataset*—form which the algorithm is learning and the method the algorithm follows to learn. Going back to our robot example, its object-to-class assignment algorithm could learn by looking at example images, classifying its objects, and comparing its decision to some *groundtruth* reference decision–much like you solving a homework problem and comparing your answer to some model answer! In this course, we will be using the word "training" as a synonym to learning experience although there is a difference between the two.

3 Paradigms of Learning Experience

There are different paradigms of learning experience (learning or training paradigms for short) which a learning algorithm may undergo. These paradigms differ in how the dataset is structured, how the performance metric is formulated, and what procedure the algorithm follows to learn. We will define four experience paradigms below.

- **Supervised learning:** This is the most common and prevalent training (experience) paradigm, especially in modern-day deep learning. Each data point in the training

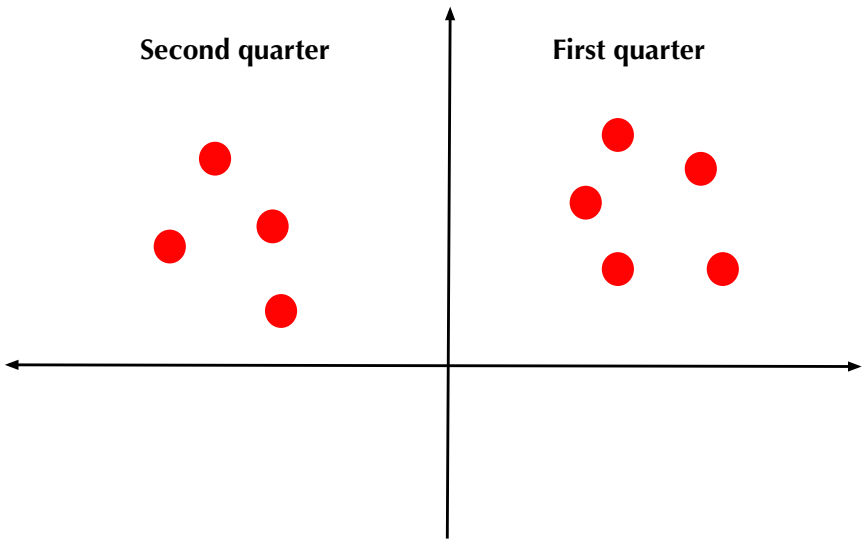

Figure 1.2: An example of unsupervised learning where an algorithm clusters data points based on common patterns.

dataset consists of an input example and a groundtuth response (usually called a target in regression and a label in classification, more on that later). The training in this paradigm is simple and very intuitive; the algorithm sees an input, produces a response, and gets assessed by the performance metric and the groundtruth response.

- **Unsupervised learning:** Opposite to supervised learning, training an algorithm here does not require groundtruth responses. The algorithm learns to discover common patterns in the input data without any form of guidance. This may seem a little too abstract, so let's consider the example in Figure 1.2. An algorithm needs to divide the set of points in the x-y plan into two groups. These points are not associated with any groundtruth responses defining their group membership, so the algorithm needs to discover a common pattern (or several patterns) to help distinguish the two groups. One clear pattern is the quarter where the points lie. If the algorithm discovers that, its job is done.

- **Reinforcement learning:** This is a learning paradigm where the learning algorithm is required to interact with its environment in order to learn. The algorithm goes through a cycle of taking an action in its environment, receiving feedback from that environment, and assessing the quality of the action from the feedback. At first, one might confuse

this with supervised learning, but it is quite different. The major difference here is in the groundtruth response. The feedback in reinforcement learning does not provide any information on what the groundtruth response looks like, rather it provides the reaction of the environment to the action the algorithm has taken.

- **Self-supervised learning:** This is one of the most recent categories in the field of machine learning. It is could be seen as a mid-way paradigm between supervised and unsupervised learning. A self-supervised algorithm is not provided with groundtruth responses in the training dataset, but it is expected to produce them from the data and the prior knowledge about the task to be learned.

The categorization of learning paradigms above is quite loos and not formally defined, and there is no consensus amongst machine learning scholars, researchers, and engineers on how to categorize learning paradigms. As such, the categorization may vary based on the background and practice of the machine learning practitioner. However, the four above could be thought of as the most popular nowadays and, hence, we will stick with them here.

Since the scope of this course is modern-practices in neural networks and deep learning, we will only focus on the first paradigm, which is supervised learning. That should not hold us back from exploring others, especially unsupervised learning, but we will do this only if time permits.

4 REMARK

At the end of this chapter, you may feel confused or a bit lost!! Do not worry; this is normal and is actually expected. So far, we have been exploring basic concepts and providing intuitive and not rigorous definitions. We will gradually go deeper into some of those concepts, and things will start falling into place for you. Just be patient and keep coming back to this chapter as we move along.

Bibliography

[1] S. Norvig, P Russel, *Artificial Intelligence: A modern approach.* Prentice Hall Upper Saddle River, NJ, USA:, 2002.

[2] I. J. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning.* Cambridge, MA, USA: MIT Press, 2016, http://www.deeplearningbook.org.