

DEMYSTIFYING SUPERVISED LEARNING A BEGINNER'S GUIDE



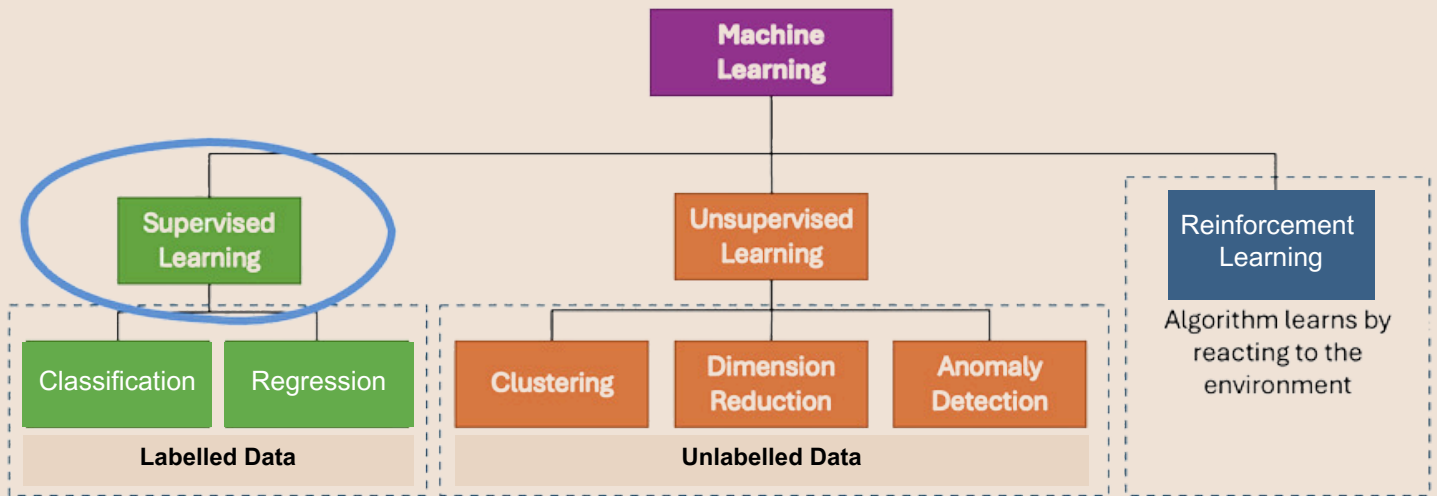
Dr Natarajan Prabhu
Senior Lecturer
NUS School of Computing

Machine Learning (ML) is a crucial subset of Artificial Intelligence (AI) that enables computers to learn patterns and insights from data. Within the realm of ML, there are three primary types of algorithms: supervised, unsupervised, and reinforcement learning.

In this article, we will delve into the intricacies of supervised learning, shedding light on its principles and applications.



Types of ML Algorithms



The goal of the supervised learning algorithms is to learn the patterns in the data and build a general set of rules to map the input to the class or events (i.e., output). Human labelled data is passed as an input to the algorithm. By analysing these labelled examples, the algorithm identifies patterns and relationships within the data, allowing it to make accurate predictions on unseen or new data.

To draw a parallel with human learning, consider how a two-year-old child distinguishes between apples and oranges. The child may rely on various features such as colour, shape, texture, taste, and smell to make this distinction. Similarly, supervised learning algorithms extract and utilise different features from the data to make accurate predictions or classifications.

“

Machine Learning (ML) is a crucial subset of Artificial Intelligence (AI) that enables computers to learn patterns and insights from data.

”

There are two primary tasks that supervised learning algorithms perform:

1. Classification

This involves categorising data into predefined classes or categories. For example, classifying an email as either spam or not spam.

2. Regression

This task involves predicting numerical values based on input data. For instance, predicting the price of a house.

In classification tasks, the algorithm receives labelled data input and produces categorical outputs such as "spam" or "not spam," "apple" or "orange," "small," "medium," or "large," etc. Conversely, in regression tasks, the algorithm takes labelled data as input and outputs numerical values, such as the price of a house, forecasted sales of a company, weather parameters, etc.

Building a supervised model involves three steps: training, testing, and classification/prediction. During training, the model learns patterns from labelled data. Testing evaluates its performance on a separate dataset that is different from the data used in training. Finally, the model is deployed to classify or predict outcomes in new unseen data instances.

Acknowledgement: ChatGPT 3.5 was used to polish this article.

**Check out our short course,
Business Applications Relying on Supervised Learning**

<https://ace.nus.edu.sg/course/business-applications-relying-on-supervised-learning/>