“Supervised Learning” Module Transcript

Chapter 1

Intro, Topics Covered, & Learning Outcomes

Hi all. My name is Hayley Bresina, and I'm a client enablement specialist on the One AI team here at One Model. In the "Introduction to Machine Learning" module, you learned that supervised learning is one of the main types of machine learning. In this module, we are going to explore supervised learning in more detail because One AI exclusively utilizes supervised learning techniques in the machine learning model tool.

We will cover an overview of supervised learning, two types of supervised learning - classification and regression, how supervised learning works at a high level, and the strengths and weaknesses of supervised learning.

After watching, you will gain a clear understanding of supervised learning as one of the main types of machine learning and its primary distinction from unsupervised learning. You will be introduced to the two main types of supervised learning, classification and regression, and begin to understand their real world applications. You will familiarize yourself with the general process of supervised learning, laying the foundation for understanding how machine learning in general works. And you will understand its strengths and weaknesses, helping you decide when it's suitable to use and when additional or different tools are more appropriate.

Chapter 2

Supervised Learning Overview

Section 2 - Supervised Learning Overview

Just a quick refresher about machine learning before diving in. Machine learning models use algorithms to learn from data, identify patterns, make predictions, or perform tasks without explicit programming. An algorithm is a set of mathematical rules or techniques that the model follows to do so.

Now we can get into our topic. Supervised learning is a type of machine learning that uses labeled datasets to train algorithms, to predict outcomes and recognize patterns, distinguishing it from unsupervised learning. While unsupervised learning also uses algorithms, it focuses on finding patterns or structures in unlabeled data. The goal of

supervised learning is to understand the relationships between input data and corresponding outputs, enabling the algorithm to make accurate predictions or decisions when presented with new unseen data. In other words, it aims to build a predictive model that can generalize by leveraging existing knowledge, which is the labeled data, to make accurate predictions or decisions about unseen data in various domains and applications.

Beyond predictions, supervised learning can be used for many purposes.

For one, categorization. In classification tasks, the goal is to assign input data points to predefined categories or classes, such as classifying emails as spam or non spam, or categorizing a group of employees into high performers and low performers.

With pattern recognition, supervised learning algorithms are used to recognize patterns and relationships in data, which can be valuable for tasks such as fraud detection or proactively addressing factors contributing to voluntary attrition.

Additionally, supervised learning helps automate decision making processes by providing recommendations based on input data. For example, in talent acquisition, it can help streamline hiring by identifying top candidates more effectively, leading to better hiring decisions and resource savings.

And finally, understanding relationships by learning from labeled data, supervised learning algorithms can uncover complex relationships between input features and output labels, providing insights into the underlying data mechanisms, such as understanding what features drive certain behaviors or how features correlate with each other.

Chapter 3

Two Types of Supervised Learning (Classification & Regression)

Section 3 - Two Types of Supervised Learning

Supervised learning can be further categorized into two main types, classification and regression. Classifications aim to predict a predefined discrete label or category. The model learns a decision boundary that separates different classes or categories in the input feature space.

For example, predicting whether an employee will voluntarily terminate in the next year or not based on features such as age, team size, manager's performance, location, and salary.

It sorts raw data into the correct categories based on patterns and characteristics it learned from the training data.

Onto the other main type. Regressions aim to predict a continuous output based on learnings from the training data.

The model learns from a mapping from input features to a continuous target variable.

Regressions help us understand the relationship between two or more variables. For example, predicting an employee's salary based on features such as tenure, education, skills, and performance ratings.

Basically, they try to predict a value based on patterns learned from training data.

The main thing I want you to take away from this slide is that classifications make predictions by sorting instances into predefined categories, while regressions predict continuous numerical values.

Both types of models are available to build in One AI depending on your problem domain and desired results. We will explore both concepts in greater detail in upcoming modules.

Chapter 4

How Supervised Learning Works

Section 4 - How Supervised Learning Works

Now we will go over the general process of how supervised learning works. We won't get too detailed as that will come in later modules, but we will aim for a basic understanding.

I will use this diagram to guide you through.

Step one is data collection and retrieval. We need datasets containing data on the features we want the model to analyze and learn from, as well as each instance's label.

An instance is the single observation or data point within the dataset.

The label is the correct answer for what category that instance belongs to. For example, if the model's job was to predict if employees will voluntarily terminate in the next year, the instances would be the employees.

Features could include salary, date of birth, manager, performance, etcetera, and the labels would be 'voluntary termination' or 'not voluntary termination'.

The model would then sort each employee into the voluntary termination label or the not voluntary termination label.

Step two is data preprocessing, where we clean the data, handle missing values, remove noise, and transform the data into a suitable format so that it can be consumed by the machine learning algorithm.

The data is also split into a training set, validation set, and test set so that we can eventually check the performance of the model.

Step three is feature extraction and selection, which is where the model chooses the features with the best predictive power for the task at hand.

Step four is choosing the strongest machine learning algorithm.

Typically, a few are tried based on the task at hand and characteristics of the dataset, and then the one that results in the best performance and fit for the model will be selected.

Step five is modeling, where the model is trained by feeding the training data into the machine learning algorithm.

This training is used to adjust the model's parameters to minimize the difference between the predicted outputs and the actual labels in the training set.

This process is often iterative with the algorithm making predictions, calculating errors, and updating the model's parameters until it's performing as well as it can.

Step six is model evaluation and tuning, which is done with the validation set of data that was set aside when the data was split. We analyze the performance of the model and refine to improve the outcomes as we see fit.

And finally, step seven is deployment and monitoring. Once we have a satisfactory model, we can deploy it to make predictions on new data in real world applications. We must monitor it over time and update it as needed to maintain effectiveness.

It's important to understand that model building is not a step by step linear process. It's very iterative, and we often go back and forward steps as needed. A model that is in use is never done. We should always be iterating.

Chapter 5

Strengths & Weaknesses of Supervised Learning

Section 5 - Strengths & Weaknesses of Supervised Learning

It's important to understand the strengths and weaknesses of any machine learning technique to know where they are applicable and where additional tools or different solutions are needed.

Let's start with the strengths. Supervised learning has a very clear objective, making it suitable for a wide range of practical applications and to various domains, including people analytics, finance, workforce planning, and talent acquisition. It can be used for tasks ranging from spam detection to quality of higher prediction.

Most supervised learning algorithms, such as decision trees or logistic regression, provide interpretable models, making it easier to understand the factors driving predictions or classifications.

This fosters trust and confidence between model builders and stakeholders because understanding the model makes translating its insights into good actionable decisions possible.

Supervised learning algorithms can achieve high performance when trained on large, high quality labeled datasets. They can learn complex patterns and relationships in the data leading to more accurate predictions.

And finally, some supervised learning algorithms support incremental learning where the models can be updated continuously with new labeled data without retraining from scratch each time. This allows the model to adapt to changing environments and evolving data distributions over time.

Moving into the weaknesses of supervised learning.

While not necessarily a weakness exclusive to supervised learning, supervised learning models are only as good as the data they were trained on. If the training data is outdated, irrelevant, contains gaps or errors, or resulted from bad or biased practices, the models created from it will inherit these issues. We can conduct bias detection and debiasing, but it's very important to remember this reality with all machine learning models.

Supervised learning also requires labeled data for training. This can be time and resource intensive to obtain and often requires domain experts to label the data correctly. Fortunately, your data in one model is already labeled.

Supervised learning may not perform well with unstructured data types, structured data types such as text, audio, and images.

And finally, some supervised learning algorithms require large amounts of labeled data to perform well. With small datasets, these algorithms may suffer from poor generalization and high variance.

Chapter 6

Conclusion & Thanks

In conclusion, supervised Learning is a powerful approach in machine learning with applicability across various domains, including people analytics.

By leveraging labeled data, supervised learning algorithms can effectively predict outcomes, classify data into predefined categories, and uncover patterns and relationships. In our next module, we will switch gears and take a closer look at unsupervised learning. Happy modeling!