“Classifications” Module Transcript

Chapter 1

Intro, Topics Covered, & Learning Outcomes

Hi, all. My name is Hayley, and I'm on the One AI team here at One Model. In the "Supervised Learning" module, you learn that supervised learning can be categorized into classifications and regressions. In this module, we are going to begin to explore classifications, which is the problem type of many of our popular One AI machine learning models, including the widely-used voluntary attrition risk recipe.

We will cover an introduction to classifications in machine learning, common classification algorithms, strengths and weaknesses of classifications, and some people analytics use cases where classifications may be applicable. After watching, you will understand the concept of classifications and their application to your organization's problem domains; be introduced to common classification algorithms, such as logistic regression, decision trees, random forest, and GBM; identify the strengths and weaknesses of classification models and learn how One AI leverages the strengths and mitigates the weaknesses; and explore practical use cases in people analytics and discover how these models can provide valuable insights to your organization.

Chapter 2

Introduction to Classifications

Section 2 - Introduction to Classifications.

A classification model is a predictive model that is trained to categorize data into distinct predefined categories based on learning from past observations. Its goal is to accurately assign the correct class label - which is the category it predicts the data will belong - to new unseen instances. The model is trained on labeled historical data where the predefined category of each instance is known. It learns to accurately class new instances by analyzing the relationships and patterns between features and class labels in the historical data.

This is how the model builds its decision function. When exposed to new data, it uses this learned relationship to predict the most likely category for each instance. The output is a categorical prediction classifying each instance.

As you can see in this image, this class is using labeled images of animals from the historical training data to learn to predict the correct label for unseen unlabeled animals from the test data.

Let's go through a People Analytics example to make this theoretical concept a bit more concrete. A common classification model in One AI is a binary voluntary attrition model. This model predicts if your current employees will voluntarily leave the organization or not in a set amount of time, usually a year.

The class labels or categories in this example are 'voluntarily terminate' and 'did not voluntarily terminate', making it a binary model since there are only two possible outcomes.

This model is trained on historical data from your organization that includes all of the relevant employee features that we want the model to learn from. With this training dataset, each employee is also labeled as either voluntarily terminated or not. The individual employee represents the instance or single observation within the dataset.

The model learns which features are associated with voluntary termination by analyzing instances where employees terminated in the training data and which features that they all shared in common. Examples of such features include lower than average salary, higher than average team terminations, and a lower performing manager.

Similarly, the model learns which features are associated with employees not voluntarily terminating by analyzing instances where employees stayed. These features could include a higher than average salary, a high performing team, and high engagement survey scores.

Using what it learned from the training data, the model then classifies current employees into the most likely category based on their features - voluntarily terminate or not - over a specified time span such as the next year. In One AI, this process occurs in a single model run where the model learns from historical data and makes predictions on the headcount population selected by the model builder.

Multi-class classifications work similarly to binary classifications in that each instance is classified into one label by the predictive model. The difference is that there are more than two possible categorical outcomes. For example, you might want to predict an employee's performance rating from high, medium, or low; so three possible categorical outcomes for a given cycle. This would require a multi-class classification.

There are also multi-label classifications where each instance can have more than one class label assigned. For example, a movie can belong to multiple genres. We won't examine this type further as One AI does not use multi label classifications.

Chapter 3

Classification Algorithms

Section 3 - Classification Algorithms

Machine learning models use algorithms to learn from data, identify patterns, make predictions, or perform tasks without explicit programming.

An algorithm is the mathematical techniques or set of rules that the model follows to do so. Classification algorithms all fall into the supervised learning algorithms category in which algorithms learn from labeled training data with known outputs.

Some common examples are the logistic regression, which despite its name, is a classification algorithm that predicts the probability of an instance belonging to a particular class using the logistic function.

It maps any numerical input into a value between zero and one, representing the probability of belonging to the positive class.

A decision tree is a tree-like model where each node represents a decision based on a feature leading to a predicted class label at the leaf nodes. It recursively splits the data into smaller parts to create subsets that are as uniform as possible in terms of class labels, continuing until a final classification is reached.

The random forest classifier is an ensemble method based on decision trees. It's like a team of decision trees working together. Each tree in the random forest is trained on a subset of the training data and a random subset of features and then makes its own prediction.

Then the forest combines all these predictions to decide the final results.

Gradient boosting machines, also called GBM, is an ensemble learning method that builds trees sequentially with each new tree correcting errors made by the previous ones. It does this by focusing on the examples that the previous trees got wrong or struggled with, essentially learning from their mistakes.

Other classification algorithms include Support Vector Machines, K-Nearest Neighbors, Naive Bayes, and Neural Networks.

Chapter 4

Strengths & Weaknesses of Classifications

Section 4 - Strengths and Weaknesses of Classifications

Like any machine learning technique, classifications have their own advantages and disadvantages.

Understanding these helps determine when they are the right tool for your project or problem domain. Let's start with the strengths. Many classification algorithms are easy to understand, making them useful for gaining insights into the relationships between input features and class labels. This interpretability helps non-experts grasp the findings and iterate on them.

Additionally, most classification algorithms are computationally efficient and scalable, handling large, high dimensional datasets with ease. This makes them suitable for real world applications, especially in people analytics where they can process these complex datasets and highlight important information, saving you time and resources.

Finally, classifications are incredibly versatile with the ability to be applied to a wide range of problems across various domains.

Beyond people analytics and One AI, they are used in fraud detection, medical diagnosis, text classifications, and more.

Moving into the weaknesses. While One AI mitigates many common weaknesses of classification models, it's still important to understand these issues, especially when using other classification systems.

Classifications can be prone to overfitting, which results in poor generalization and performance on unseen data because they learn the training data a bit too well. One AI reduces this risk with cross validation and other techniques.

Additionally, if the training dataset is imbalanced, meaning most instances fall into one category, it can lead to bias models that favor the majority class and perform poorly on minority classes. One AI addresses this with upsampling techniques to balance the datasets. And finally, many classifications out in the wild are black boxes, meaning their internal workings are not interpretable, making it difficult to understand how it arrived at its predictions or decisions.

One AI provides transparency with exploratory data analysis (EDA) and results summary reports for each run with everything you need to know about how the model was configured and performed and made its feature selections.

Chapter 5

People Analytics Use Cases

Section 5 - People Analytics Use Cases where Classifications are Applicable

Classification models are powerful tools in people analytics, enabling organizations to extract valuable insights from employee and talent acquisition data.

By categorizing employees based on attributes such as performance, engagement, turnover risk, or skill sets, these models help in making informed decisions about talent management, workforce planning, employee development, and retention strategies.

They identify patterns, trends, and predictors of employee behavior leading to more effective human resources practices and improved organizational performance.

Often the insights into the drivers of these behaviors are actually more valuable than the predictions themselves.

Some common classification models that you can build in One AI include voluntary attrition, which predicts whether an employee will voluntarily leave the organization after a set period, typically one year, Involuntary attrition, which is just like the voluntary attrition one, but targets involuntary terminations. New hire failure and success, which predicts if new hires will remain with the organization after a set period; again, typically one year. We also have promotion models, which predict if an employee will be promoted within a set period of time, and high performer models, which predict if an employee will achieve high performance ratings within a set period.

Each model also provides detailed insights into what is driving these behaviors to get you a more robust understanding of your workforce and their motivations.

This helps with the development of strategies to address problem areas.

Additionally, One AI allows for custom classification models using any metric in one model with our guided framework. For example, you can predict if an employee will need a performance improvement plan or if they will be highly engaged over the next year. As long as you have the appropriate data in One Model to train the model and a defined metric for what you want to predict, the sky's the limit.

Chapter 6

Conclusion & Thanks

In conclusion, classification models are powerful tools in machine learning and people analytics. They categorize data into distinct classes based on past observations, enabling organizations to extract valuable insights and make informed decisions about talent management, workforce planning, and employee retention.

By taking what you've learned in this module, you can effectively leverage classifications to drive positive outcomes in your organization. Next, we will explore the other half of supervised learning - regressions. Happy modeling!