“Intro to ML” Module Transcript

Chapter 1

Intro, Topics Covered, & Learning Outcomes

Hi, everyone. My name is Taylor Clark. I'm the Chief Data Scientist and Head of Machine Learning here at One Model. And I am here to talk to you about machine learning, what it is, how you can use it, and a number of other topics.

So, we'll cover a few things today. We're going to look at what exactly machine learning is, generally, how machine learning works. We'll cover some strengths and limitations, which is a particularly useful topic, and then we'll hit on some real world applications of machine learning.

For learning outcomes, we're really hoping that you'll understand kind of how machine learning works from a broad level, where it can be used, where it might be less applicable, the importance of building models correctly and spending time understanding, what drives them as well as, some of the fundamental terminologies and steps required to work through machine learning.

Chapter 2

Machine Learning Example

All right. So, before we kind of get into specific terminology, we'll cover some intuition for what machine learning is. So we have a graph in front of us right now, and this is some different ratings about movies that I either liked or disliked, analyzing them based on run time, either short to long, or comedy level, highly funny or kind of low levels of comedy. And you can see the green dots represent films that I enjoyed, and the orange dots represent films that I didn't enjoy.

So there's a kind of intuition that we can see here. These are, you know, historical data points with labels on them, for whether I dislike them or not. But, generally, the longer the film is, the more likely I am to enjoy it. The higher the comedy level, the more likely I am to enjoy it.

And then the inverse is true. The shorter the film and the less humor in the film, I'm kind of more likely to not enjoy the film.

So, we can look at this, and as humans, we can interpret and learn from this. If we applied an algorithm, through machine learning, we could look at this data and we could decide we could kind of suss out some of those same characteristics and say, okay., generally, longer films, you know, we enjoy more. Shorter films, we enjoy less, and make those observations.

So, once we've learned that as humans or as a machine, we can insert new pieces of data. So if we put an orange dot here, this big orange dot in the top right, that indicates a new film that we haven't seen yet, and we could make a pretty good prediction that, hey, it's it's likely that we'll enjoy this film given, it's high in comedy and it's generally longer. We don't know that with a certainty, but, you know, based on those two characteristics, it's probably a good bet.

The blue dot that we added to the chart here is more ambiguous. We're not certain how to classify this in terms of whether I would like it or dislike it. It is kind of right in the middle, so maybe there's a probability 50% chance that I'll enjoy it, 50% chance that I won't. I would guess just given the fact that it's a little bit lower on the comedy level from a personal level, I probably won't enjoy it as much. So, this is an example of how we as humans can apply intuition to a chart to make predictions, and then we're going to try to use machine learning to accomplish the same types of things.

Chapter 3

Machine Learning Overview

So, there's a lot of terminology in machine learning. It can generally be a little bit confusing.

So, we'll try to hit a few common terms here.

Machine learning is just the use of algorithms to look at and learn from data to make observations on the data and or make predictions into the future or into data that we haven't seen before.

So, on the right, I like this diagram a lot. Artificial intelligence is used really broadly in the current market, and there's maybe not as much understanding about what that specifically means. Generally, it's going to be some sort of algorithm that people are using to emulate human performance on certain tasks. And it's very similar fashion to what we did in the previous chart where we looked at it with our human intuition and we made predictions. But instead of using my eyes to do it, we'll use algorithms to do it.

Within artificial intelligence, there's a lot of different techniques, reinforcement learning, supervised learning, unsupervised learning.

So, supervised learning is generally going to be very similar to what we just looked at where we have pre labeled data that we're training from. So we're looking at, you know, liked or disliked movies, and we're learning from that labeled data because we know that I liked it or didn't. And then we're kind of making predictions into films that we haven't seen before. Unsupervised would be when the table's data is unlabeled.

So that might be building clusters off certain types of personas, in your client data or user data where we don't have predefined categories, but instead we're using algorithms to help us detect where good categories might exist.

And then reinforcement learning is generally just going to integrate some feedback within the process to get better and better over time.

Chapter 4

How Machine Learning Works

Okay. Now we're going to go into a bit more detail on how machine learning actually works.

I like this flow diagram. So, we're going to focus on the broader process of machine learning, not necessarily the specifics of the algorithms. So, generally, you're gonna have some set of input data. This could be the example that we're hitting earlier around films that we liked or didn't like. It could be records about employees and whether they terminated or not or they had successful promotions or not. And we're going to do some set of preparation on that data. So that might be, looking at it for duplicates, applying some statistical transformations to represent data in different ways, as well as just general data cleaning.

The next step is what we like to call exploratory data analysis. So, this is a really critical step that often gets overlooked and can be immensely valuable prior to actually doing models. This is where you might go and look at descriptive statistics about the relationships between the training datas and the outcomes that you're looking at. So the visual effort that we did, looking at films that I liked or didn't like, we were able to identify certain relationships. And so exploratory data analysis would be a more technical application of that where we go and we look at correlations. We look at where, you know, empty data is occurring. Just any characteristics that might be insightful or useful to us, in the data prior to actually trying to fit an algorithm to the data to help us make predictions to an unforeseen data.

So once we understand the data, and we've prepared it, we'll move into the actual model training pieces. This is where either we're going to use a set of different algorithms, or we'll use a single algorithm with certain settings, and we will apply that algorithm to the data so that it can learn from the data either in a reinforcement manner or in a supervised manner or an unsupervised manner. In the supervised case with the movie films, we're going to look at the labels, which would be whether I like the film or not and the characteristics about the film, then to make predictions. And there's a lot of different algorithms out there, and that process is generally called training or fitting.

Once we've trained and or fitted a model to the data, then we can apply it to new data. This could be data in the future that we're trying to make predictions into the future or data that we just haven't seen with the algorithm. So films that we don't know whether I'll like or not.

And we can make predictions against that data, say, yeah. We, you know, we think Taylor's likely to enjoy this film or not. And then from that, we're going to hopefully be making some decisions, kind of incorporating both the data analysis and the predictions from model and some data about how well the model performed to help us decide, in this case, what movies we might want to watch. So maybe we have a list of 100 potential movies we want to watch, and I want to decide, you know, which ones I'm likely to enjoy. I'll build that algorithm, apply it to the data, and come up with a list of 5 films that I feel really confident I'm going to enjoy.

Chapter 5

Strengths & Limitations of Machine Learning

So, strengths and weaknesses of machine learning. I like to talk about machine learning as a tool. It's not the answer in and of itself, but it's a tool that we can use in our arsenal to help us make decisions. So, from a strength perspective, it can be really efficient, and it can scale really well.

So, if we had lots of different parameters outside of just movie length and comedy level that we wanted to take into consideration in our film recommendation algorithm, Could add in a bunch of those things and run the analysis in a way that might be hard for us to visualize on on just a regular graph as a human. So we got a lot more parameters, and then also that recommendation engine or algorithm that we developed, we could apply it really quickly to films without, you you know, having to go to IMBD, look it up, and decide whether I'm going to like it or not, which also allows us to kind of scale that out in an automated fashion as well.

There's a lot of flexibility in what you can do with these algorithms in terms of predictions and forecasts and classifications, and this data can be immensely valuable. And machine learning tends to be very good at those specific tasks. In terms of limitations, I think it's critical that we all talk about the limitations and acknowledge that there are many.

So, the data dependency, if you don't have a lot of training data, it's going to be hard to get an algorithm to perform well. That training data might not be representative of the data that you're trying to apply the algorithm to. Maybe all the films we're looking at are, you know, medium length, and, they're not comedies, or kind of the comedy is of of a different characteristic that we're not interested in. It might be hard for our algorithm before to do well on that because none of them are comedy, so we'll just assume that I will not like all of them, which may or may not be true.

Interpretability. Some of the algorithms can be hard to interpret, and the different techniques that you use to interpret them can have different degrees of precision. So, you know, in situations where understanding the causal mechanism for why a prediction is being made with absolute certainty, sometimes machine learning might be complicated and difficult to use in those situations.

Overfitting, concerns. Generally, the process in which you train an algorithm or machine learning is very important. It's actually quite easy to do it poorly and result in an algorithm that makes predictions, but it doesn't do so well. So understanding the process and making sure that you're able to spend the time developing a rigorous process is critical in ensuring that you'll have a quality machine learning algorithm.

Bias and fairness, also extremely important. If the data that you're using has biases in it, it's likely that the algorithm that you're using is going to perpetuate those biases, but it kind of intersecting that with interpretability. It might be hard to detect that. There are lots of bias detection algorithms out there that should be applied to machine learning, and there's metrics you can look at as well. But, again, if the data is biased, and there's some concerns about fairness and things like that, often that can propagate through.

Chapter 6

Applications of Machine Learning

Applications of machine learning

So, we'll look at a few different ways in which people use machine learning across industries. Won't spend too much time on this. I think this diagram does a good job of laying out a bunch of different use cases. But yeah. Like, in any case where you have a good amount of data and it might be useful to automate certain types of analysis and or predictions, machine learning can be very valuable.

Most of us interact with spam filters positively or negatively. Those are generally powered by machine learning algorithms, spell checking, and healthcare. There's lots of applications in health care and fraud detection more broadly.

In the HR space, you know, we see people looking at flight risk, new hire failure rates, performance analyses, things of that nature. There's lots of room for the exploratory data analysis as well as predictions throughout industries.

Chapter 7

Conclusion & Thanks

So, yeah, this is just the intro. There's lots of more detailed content to come. But thank you for watching, and I hope you enjoy your time at One AI Academy!