

Question we can have it predict future

Business, Decision Making



Question 1 Give a definition of a learning system. What are the basic components of the learning system? There are three different tasks in machine learning (ML): supervised learning, unsupervised learning, and reinforcement learning.

How are they different? What kind of data is provided for each learning task?

How is reinforcement learning (RL) related to supervised learning? How is RL related to unsupervised learning? Explain how to teach a chess—playing

program using supervised, unsupervised, and reinforcement learning

paradigms, respectively. Learning is any practice by which a system improves

Performance from experience. For example, let take school as an example.

School have teacher, books and other resources. Students they have they

get experience from their teacher, books and other resources and as a result

of this they improve their performance. Basic

components

I. Performance II.

Task (future task) III. Training experience (datasets) Supervise

means to observe and direct the execution of a task. Means supervising a machine learning model that might be able to produce classification regions.

Teach the model by training it with some data from a labeled dataset than

load the model with knowledge so that we can have it predict future

instances. Generally speaking, the model is trained on a labeled dataset, so it

can predict the outcome of out of sampled data. There is 2 type of supervised

learning: classification and regression. Unsupervised learning is exactly as it

sounds, let the model work on its own to discover information that may not

be visible for our eyes. It uses machine learning algorithms that

extract conclusions on unlabeled data.

Unsupervised learning has more difficult algorithms than supervised learning, since we know little to no information about the data, or the outcomes that are to be expected. With unsupervised learning, we're looking to find things such as group, clusters perform density estimation and dimensionality reduction. In supervised learning, however, we know what kind of data we're dealing with, since it is labelled data. A system interacts with environment to perform some task in return of this environment give some reward to the system reward is either positive or negative on the base of this reward system improve his performance and again perform task until system achieve maximum positive reward. Reinforcement learning is something between supervised and unsupervised learning it told us when we are wrong through negative reward but it didn't tell us which way to get maximum positive reward it should explore all the possibility. In comparison to supervised learning, unsupervised learning has: fewer tests and fewer model that can be used in order to ensure the outcome of the model is accurate. As such unsupervised learning create a less controllable environment, as the machine is creating outcomes for us.

The biggest difference between supervised and unsupervised learning is that supervised learning deals with labeled data while unsupervised learning deals with unlabeled data. In supervised learning, we have machine learning algorithms for classification, and regression. Classification is the organization of labeled data and regression is the prediction of trends in labeled data to determine future outcomes. In unsupervised learning, we have clustering. Clustering is the analysis of patterns and groupings of unlabeled data. Reinforcement learning is not exactly supervised learning because it doesn't

rely strictly on labeled data. It actually relies on “reward”. But it’s not unsupervised learning either, since we know straight when we model our “learner” which is the estimated reward.

Supervised learning: Labeled Data. Unsupervised learning: Unlabeled Data. Reinforcement learning: Have no data we construct a model that generates data based on reward.

Question 2 Consider solving the problem of character—image classification using a back-propagation learning neural network. What is the model structure? What training data should be provided? How does the neural network (model) learn from data? Explain the procedure for learning.

Question 3 What is generalization? We’ve mostly been talking about the training data right so and classifiers and the algorithms use the training data to build the predictors so the training data attributes X along with some targets Y which could be a class or a number or something so that’s what we used to train to build our model but then the reason we’re building this predictor is that sometime tomorrow or in the future we’re going to get new data and on that data.

What is overfitting in learning? Overfitting is happened when you find or when you learn a predictor that fits the training data a little bit too well so it’s usually happens when your predictor the function that you’re predicting is complex enough and flexible enough to fit any kind of sort of noise in the training data so those are patterns that are present in the training data that will not be present tomorrow in the future data that you see so the when that happens you say you over fit. Explain the overfitting phenomena using the

polynomial interpolation problem. How can we avoid the overfitting problem in learning? Debugging and diagnosing things they can go wrong with learning algorithms will give you specific tool to recognize when overfitting. If we think overfitting is occurring what can we do to address when we had 1 or 2 dimensional data so we could just plot the hypothesis and see what was going on and select the appropriate degree polynomial. We could just plot hypothesis and if it was fitting the sort of very wiggly function that goes all over the place to predict and we could then use appropriate degree polynomial. So plotting the hypothesis could be one way to try to decide what degree polynomial to use but that does not always work. In fact it when we have so many features it also becomes much harder to plot the data and becomes much harder to visualize it to decide what features to keep or not so concretely suppose if we're trying to predict housing prices sometimes we can just have a lot of different features and all of these features seem you know maybe they seem kind of useful but if we have a lot of features and very little training data then overfitting can become a problem in order to address overfitting there are two main options for things that we can do the first option is to try to reduce the number of features concretely one thing we could do is manually look through the list of features and use that to try to decide which are the more important features and therefore which are the features we should keep and which other features we should throw out there are algorithms for automatically deciding which features the key and which features to throw out this idea of reducing the number of features can work well and can reduce overfitting and when we talk about model selection we'll go into this in much greater depth but the disadvantage is that by throwing

away some of the features is also throwing away some of the information you have about the problem for example maybe all of those features are actually useful for predicting the price of a house so maybe we don't actually want to throw some of our information or throw some of our features away.

Regularization: we're going to keep all the features but we're going to reduce the magnitude this method works well we'll see when we have a lot of features each of which contributes a little bit to predicting the value. How can we avoid the overfitting problem in back-propagation neural networks?

Downhill If you have very uneven training and test data structure, try to fix it.

E. g.

the share of classes zero and one in both datasets should be equal. You can also randomly drop some neurons during training. For randomly stopping neurons you can use the L2 loss function

Question 4 Give a list of

machine learning models for supervised learning. How are they different?

What are their similarities? What representations do they use? What ML

methods use tree structures for representing their model? What methods use

graph or network representations? What methods use list structures or rule

sets for representing models? o Support Vector

Machines o linear regression o logistic regression o naive Bayes o linear discriminant analysis o decision trees o k-nearest neighbor

algorithm o Neural Networks (Multilayer perceptron). Classification (1R, Naive Bayes, Decision tree learning algorithm such as ID3 CART and so on) Numeric

Value Prediction Decision tree continuous and categorical inputs. While decision trees classify quickly, the time for building a tree may be higher than

another type of classifier Decision trees suffer from a problem of errors propagating throughout a tree Decision trees can be used to help predict the future The trees are easy to understand Decision trees work more efficiently with discrete attributes The trees may suffer from error

propagation SVM continuous value inputs Naïve Bayes A

simple but effective learning system. Each piece of data that is to be classified consists of a set of attributes, each of which can take on a number of possible values. The data are then classified into a single classification.

Advantages: -Fast to train (single scan). Fast to classify -Not sensitive to irrelevant features -Handles real and discrete data -Handles streaming data well Disadvantages: -Assumes independence of features KNN

continuous value inputs

- o Decision Trees are fast to train and easy to evaluate and interrupt.
- o Support vector machine gives good accuracy, power of flexibility from kernels.
- o Neural network is slow to converge and hard to set parameters but if done with care it works

well so Bayesian classifiers are easy to understand. Question

5 Machine learning methods can be defined by three dimensions: type of learning data, model structure, and learning algorithm. Describe these aspects for each of the following methods.

Linear regression Type of learning data Continuous and categorical inputs Model structure Linear regression is a very simple approach for supervised learning. Though it may seem somewhat dull compared to some of the more modern algorithms, linear regression is still a useful and widely used statistical learning method. Linear regression is used to predict a

quantitative response Y from the predictor variable X . Linear Regression is made with an assumption that there's a linear relationship between X and Y . $Y = W_0 + W_1X$, where X is the explanatory variable and Y is the dependent variable. The slope of the line is W_1 , and W_0 is the intercept (the value of y when $x = 0$).

Learning algorithm: Multi-layer perceptrons, Support vector machine, Random forest
 Decision trees: Type of learning data: Continuous and categorical inputs
 Model structure: A tree has many analogies in real life, and turns out that it has influenced a wide area of machine learning, covering both classification and regression. In decision analysis, a decision tree can be used to visually and explicitly represent decisions and decision making.

As the name goes, it uses a tree-like model of decisions. Though a commonly used tool in data mining for deriving a strategy to reach a particular goal, it's also widely used in machine learning, which will be the main focus of this article. The goal of Decision Tree is to create a model that predicts the value of a target variable by learning simple decision rules inferred from the data features.

Learning algorithm: Random forest, C4.5, ID3, C5.0 etc.
 Neural networks: Type of learning data: Continuous and categorical inputs
 Model structure: The dendrites carry the signal to the cell body where they all get summed. If the final sum is above a certain threshold, the neuron can fire, sending a spike along its axon. In the computational model, we assume that the precise timings of the spikes do not matter, and that only the frequency of the firing communicates information.

we model the firing rate of the neuron with an activation function (e. g. sigmoid function). Learning algorithm K-means clustering Type of learning data Continuous inputs Model structure K-means clustering is a type of unsupervised learning, which is used when you have unlabeled data (i. e., data without defined categories or groups). The goal of this algorithm is to find groups in the data, with the number of groups represented by the variable K. The algorithm works iteratively to assign each data point to one of K groups based on the features that are provided.

Data points are clustered based on feature similarity. Rather than defining groups before looking at the data, clustering allows you to find and analyze the groups that have formed organically. The “ Choosing K ” section below describes how the number of groups can be determined. Each centroid of a cluster is a collection of feature values which define the resulting groups.

Examining the centroid feature weights can be used to qualitatively interpret what kind of group each cluster represents. Pizza example which you told us in the class population divided according to 3 branches of pizza.

Learning algorithm Naive Bayes Classifier Type of learning data Model structure This lets us examine the probability of an event based on the prior knowledge of any event that related to the former event. So for example, the probability that price of a house is high, can be better assessed if we know the facilities around it, compared to the assessment made without the knowledge of location of the house.

Bayes' theorem does exactly that. Above equation gives the basic representation of the Bayes' theorem. Here A and B are two events and, $P(A|$

$P(A|B)$: the conditional probability that event A occurs, given that B has occurred. This is also known as the posterior probability. $P(A)$ and $P(B)$: probability of A and B without regard of each other.

$P(B|A)$: the conditional probability that event B occurs, given that A has occurred. Learning algorithm Question 6 Explain how machine learning can be used for the following applications. You may refer to existing work and discuss it. Optical character recognition Predicting customer's response to coupon mails Speech recognition Autonomous car driving Clustering the types of customers of internet shopping malls Board game playing (e. g., Chess, backgammon) Helicopter control