Machine Learning Technique

(Nuanwan Soonthornphisaj : 12 June 2000)

There are many unsolved problems that computers could solve if the appropriate software existed. .Many of these problems are currently unsolvable, not because current computers are too slow or have too little memory, but simply because it is too difficult to determine what the program should do. If a computer could learn to solve the problems through trial and error, that would be of great practical value.

What is machine learning ?

Machine Learning is concerned with computer programs that automatically improve their performance through experience. Machine Learning methods have been applied to problems such as learning to drive an autonomous vehicle, learning to recognize human speech, learning to detect credit card fraud, and learning strategies for game playing.

- ## Supervised Learning

For a system to be using supervised learning, a teacher must help the system in its model construction by defining classes and providing positive and negative examples of objects belonging to these classes. The system is then to find out common properties of the different classes, and what separates them, in order to make correct classification for other objects.

Supervised learning is a general method for training a parameterized function approximator, such as a neural network, to represent functions. However, supervised learning requires sample input-output pairs from the function to be learned. In other words, supervised learning requires a set of questions with the right answers. For example, we might not know the best way to program a computer to recognize an infrared picture of a tank, but we do have a large collection of infrared pictures, and we do know whether each picture contains a tank or not. Supervised learning could look at all the examples with answers, and learn how to recognize tanks in general.

- ## Unsupervised Learning

Collecting and labeling a large set of sample patterns can be very expensive. For example recorded speech is virtually free, but accurately labeling the speech can be very expensive and time consuming. By designing a basic classifier with a small set of labeled samples, and then tuning the classifier up by allowing it to run without supervision on a large, unlabeled set, much time and trouble can be saved.

Training with large amounts of often less expensive, unlabeled data, and then using supervision to label the groupings found. This may be used for large "data mining" applications where the contents of a large database are not known beforehand. Characteristics of patterns can change slowly with time. For example in automated food classification as the seasons change. If these changes can be tracked by a classifier running in an unsupervised mode, improved performance can be achieved.

In the case of *unsupervised learning*, no teacher defines the classes a priori. Thus, the system itself must find some way of clustering the objects into classes, and also find descriptions for these classes. The resulting rules from such a system will be a summary of some properties of the objects in the database: which classes are present and what discerns them. This will of course only be what the system has found as most prominent, but there may be many other ways of dividing the objects into classes, and many ways of describing each class.

Unsupervised learning methods require only the input and determine by themselves which classes exist in the input and which possible inputs belong to each class.

- Reinforcement Learning

Unfortunately, there are many situations where we don't know the correct answers that supervised learning requires. For example, in a flight control system, the question would be the set of all sensor readings at a given time, and the answer would be how the flight control surfaces should move during the next millisecond. Simple neural networks can't learn to fly the plane unless there is a set of known answers, so if we don't know how to build a controller in the first place, simple supervised learning won't help.

For these reasons there has been much interest recently in a different approach known as reinforcement learning (RL. Reinforcement learning combines the fields of dynamic programming and supervised learning to yield powerful machine-learning systems. Reinforcement learning appeals to many researchers because of its generality. In RL, the computer is simply given a goal to achieve. The computer then learns how to achieve that goal by trial-and-error interactions with its environment. Thus, many researchers are pursuing this form of machine intelligence and are excited about the possibility of solving problems that have been previously unsolvable.

Reference

Kaelbling, L. P., Littman, M. L., and Moore, A.W. (1996). Reinforcement Learning: A Survey, Journal of Artificial Intelligence Research Volume 4, pp. 237-285.