CPS331 Lecture: Introduction to Learning; Fogel Discussion

last revised October 20, 2016

Objectives:

1. To understand the structure of a learning program
2. To understand the basic ideas of supervised, reinforcement, and unsupervised learning.

Materials:

1. Projectable of Learning Program Structure
2. Projectable of Fogel discussion question

I. Introduction

A. One of the crucial hallmarks of intelligent life is the ability to learn. In fact, the ability to learn tricks and obedience to commands is one way we distinguish more intelligent pets (such as dogs) from less intelligent ones (such as guinea pigs.)

B. Thus, enabling computer programs to learn is one crucial area that must be addressed in any attempt to produce high levels of intelligence.

1. In one sense, learning can be regarded as a component of almost any AI problem. (Cawsey points out some examples of this.)

2. In another sense, though, it is useful to study learning as an AI problem in its own right - understanding that ultimately learning becomes a part of a larger system that solves some problem.

3. Recall what Turing called "Lady Lovelace's Objection" to the idea of an intelligent computer - a computer can only do what it is programmed to do. In a sense, learning seeks to move us beyond that (though of course one might object that a learning computer is still doing what it is programmed to do - learn!)

C. Of course, we must at some point define what we mean by "learn".

1. Tanimoto gives the following definition: "When a system learns, it improves its knowledge or its ability to perform one or more tasks. The improvement comes about as a result of information-processing activity."

   a) Note that the final phrase - "as a result of information processing activity" is intended to distinguish learning from simple programming.

   b) Example: We have already discussed how an expert system might be constructed by writing "if-then" rules. This is a form of programming. An expert system that learns, on the other hand, would infer its rules from example "cases".

2. Russell and Norvig suggest that learning can be understood as the system acquiring the ability to compute an approximation to a function - the function that maps an input situation to the correct response.

D. In general, a learning system has the following structure:

PROJECT

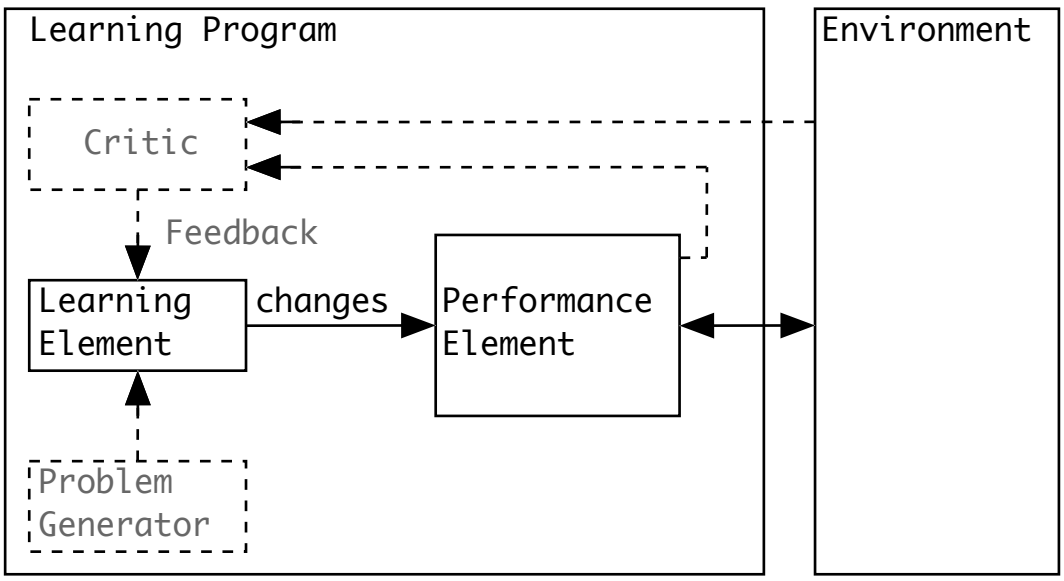

1. The performance element corresponds to a program without learning. It does what the program is ultimately intended to do.

   a) For example, if the learning program learns to play a game, then the performance element is what actually plays the game.

   b) The performance element interacts with the environment. This may be the actual environment, or a simulated one.

      For example, if the learning program learns to play a game, then the environment may be the external interface plus the opponent, or it may be another copy of the program that the learning program plays against.

2. The learning element updates the performance element.

3. The critic provides feedback on the the performance element's performance, based either on examining its output or on examining the effect of its output on the environment.

4. We will discuss the problem generator as we discuss various paradigms.

II. **Learning can occur in several ways.**

A. One approach is known as <u>supervised learning</u>. This approach is often used with classification problems, which have the following general structure

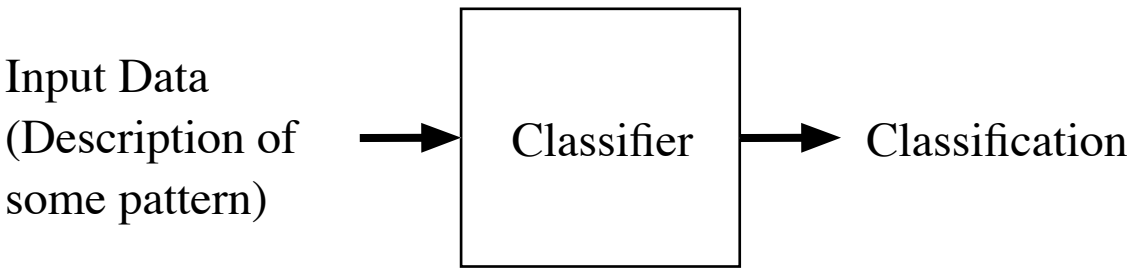

B. The classification may be as simple as “instance of the concept or not an instance of the concept”, or it may involve multiple categories.

Example: A familiar example is the “junk mail” (spam) filter that is part of many email programs.

1. There is a set of training data, which consists of pairs of the form (input data, classification), often drawn from data about past observations.

   a) The problem generator presents the pairs to the performance element one by one.

   b) The critic compares the classification produced by the performance element with the correct classification from the training data.

      (1) If the performance element gives an incorrect classification, then the learning element modifies the performance element appropriately.

      (2) The system is expected to learn how to correctly classify each training pattern, with the expectation that it will then be able to correctly classify similar patterns that are not part of the training data.

   c) Under this approach, there is a “teacher” who chooses the training data to be presented to the system. However, this differs from conventional program maintenance in that the teacher does not directly modify the knowledge base; rather, the learning element of the program modifies the performance element in response to how it handles the data presented to it.

      Example: A junk mail filter is typically trained as follows

      (1) If the user determines that an incoming piece of mail is junk, the user indicates this in some fashion (e.g. by clicking a

button). This causes the filter to adjust its classification rules so that similar mail in the future will tend to be classified as junk.

(2) A junk mail filter generally quarantines incoming mail it classifies as junk in a special "junk mail" folder. If the user determines that a piece of mail the filter has placed in this folder is not really junk mail, he/she can indicate in some way (e.g. by clicking a different button) that this particular piece of mail is not junk. This causes the filter to adjust its classification scheme so that similar mail in the future is not likely to be classified as junk.

(3) Notice, the use of words like "tend to" and "not likely" in the above. Learning is typically incremental, so it may take human classification of several pieces of mail for the filter to "learn" the correct scheme.

2. Can you think of examples where humans learn this way?

   ASK

   Example: we might use this approach to teach a child to recognize different kinds of animals. The teacher shows the child pictures of various animals and tells the child what each is. The hope is that the child will learn to recognize other similar animals when shown a picture.

C. A second approach is known as <u>reinforcement learning</u>

1. The system performs either in its intended environment or a simulated environment, with the only feedback being whether or not its performance is correct. (The system is not given the correct answer if it did the wrong thing).

   For example, if the learning program plays a game, the feedback may be simply whether or not the program won the game.

a) The feedback is sometimes referred to as a "reward" - with a positive reward if the program did the right thing, and a negative reward it if did not.

b) Whenever the critic produces a negative reward, the learning element updates the performance element appropriately. The expectation is that the program will eventually learn to make the correct choice when confronted with similar situations in the future.

2. There may or may not need to be a problem generator, per se, depending on whether the system is designed to learn in a real environment or a simulated one.

3. Again, can you think of examples where humans or animals learn this way?

   ASK

   a) Example: we learn a lot of skills this way ourselves: riding a bicycle, shooting a basketball through the hoop, playing a particular game ...

   b) Example: training a dog [reinforcement is either a treat or a scolding].

   c) Example: if you've ever watched a squirrel trying to get into a bird feeder, you know that many animals learn this way too!

D. A third approach is known as <u>unsupervised learning</u>.

1. Rather than being presented with right or wrong answers, the learning system extracts patterns from data, or conducts some sort of "thought experiment."

2. In this approach, there really is no problem generator or critic per se.

3. Among humans, the ability to learn in this way distinguishes an expert in a field from a neophyte. This approach to learning is exemplified by the theses or projects required in many undergraduate or graduate programs.

4. However, this kind of learning is not limited to human experts. In fact, for a new-born child this is the first kind of learning the child does, as he/she gradually discovers his/her bodily faculties and how to control them! For example, initially the sounds young children make are meaningless babble, but as the child begins to learn how to speak he or she begins to focus on making the sounds that the child hears others in its environment making - that is, the phonemes of the language the child is about to learn.

E. Learning systems come in several rather distinct "flavors".

1. Symbolic learning - in which the system learns rules, or parameters of rules.

2. Genetic learning - a very different approach, modeled on natural selection.

3. Connectionist learning - in which the system learns connection weights in a neural network.

4. Note that supervised, reinforcement, and unsupervised learning can all be used with either a symbolic or connectionist learning system. Genetic learning is basically a reinforcement paradigm.

5. In the next lecture, we will focus on symbolic learning approaches. We will discuss genetic and connectionist learning in subsequent lectures.

6. Later, we will focus on some more novel approaches inspired by biological systems that cannot perhaps be considered learning in the strictest sense, but are nonetheless interesting.

F. Because human learning, in general, is not well understood, it should not be surprising that machine learning is still very much a research frontier, with very different approaches used to solve different problems. The literature on this topic is extensive.

III. Discussion of Fogel Chapter

A. In general, what was your reaction to this chapter?

B. How is Fogel's "take" on AI different from that of Turing or Newell and Simon?

1. According to Fogel, how has symbolic AI defined intelligence?

2. How would Fogel define "intelligence"?

3. What is Fogel's opinion of the Turing test and "classic" AI work such as chess-playing programs like Deep Blue? Do you agree or disagree?

C. In Fogel's view, how should we go about trying to produce intelligent machines?

D. Small group discussion: Do you think Fogel's critique of symbolic AI is on target? Why or why not?

PROJECT

E. Actually, similar views have been espoused by others as well. I hope we'll hear some things along these lines from the group presenting on Situated AI / Behavior-Based Robotics, for example.