Artificial intelligence and advanced robotics techniques

The next sections will each detail a step-by-step example of the application of a different AI approach.

We start with object recognition. We need our robot to recognize objects, and then classify them as either toys to be picked up or not toys to be left alone. We will use a trainedartificial neural network (ANN) to recognize objects from a video camera from various angles and lighting conditions.

The next task, once a toy is identified, is to pick it up. Writing a general purpose pickup anything program for a robot arm is a difficult task involving a lot of higher mathematics(google inverse kinematics to see what I mean). What if we let the robot sort this out for itself? We use genetic algorithms that permit the robot to invent its own behaviors and learnto use its arm on its own.

Our robot needs to understand commands and instructions from its owner (us). We use natural language processing to not just recognize speech, but understand intent for the robot to create goals consistent to what we want it to do. We use a neat technique called the“fill in the blank” method to allow the robot to reason from the context of a command. Thisprocess is useful for a lot of robot planning tasks.

The robot's next problem is avoiding the stairs and other hazards. We will use operantconditioning to have the robot learn through positive and negative reinforcement where it issafe to move.

The robot will need to be able to find the toy box to put items away, as well as havea general framework for planning moves into the future. We will use decision trees for pathplanning, as well as discuss pruning for quickly rejecting bad plans. We will also introduceforward and backwards chaining as a means to quickly plan to reach a goal. If you imaginewhat a computer chess program algorithm must do, looking several moves ahead and scoringgood moves versus bad moves before selecting a strategy, that will give you an idea of the power of this technique. This type of decision tree has many uses and can handle many dimensions of strategies. We'll be using it to find a path to our toy box to put toys away.

Our final practical chapter brings a different set of tools not normally used in robotics, or at least not in the way we are going to employ them.

I have four wonderful, talented, and delightful grandchildren who love to come andvisit. You'll be hearing a lot about them throughout the book. The oldest grandson is six years old, and autistic, as is my grandaughter, the third child. I introduced the grandson, William, to the robot , and he immediately wanted to have a conversation with it. He askedWhat's your name? and What do you do? He was disappointed when the robot made no reply. So for the grandkids, we will be developing an engine for the robot to carry on a smallconversation. We will be creating a robot personality to interact with children. William hadone more request of this robot: he wants it to tell and respond to knock, knock jokes.

While developing a robot with actual feelings is far beyond the state of the art in robotics or AI today, we can simulate having a personality with a finite state machine and some Monte-Carlo modeling. We will also give the robot a model for human interaction sothat the robot will take into account the child's mood as well. I like to call this type of software an artificial personality to distinguish it from our artificial intelligence. AI builds amodel of thinking, and AP builds a model of emotion for our robot.