


Since few readers are likely to attempt building a learning machine that requires 300 matchboxes, I have designed hexapawn, a much simpler game that requires only 24 boxes. The game is easily analyzed—indeed, it is trivial—but the reader is urged *not* to analyze it. It is much more fun to build the machine, then learn to play the game while the machine is also learning.

Hexapawn is played on a 3×3 board, with three chess pawns on each side as shown in Figure 35.1. Dimes and pennies can be used instead of actual chess pieces. Only two types of move are allowed: (1) A pawn may advance straight forward one square to an empty square; (2) a pawn may capture an enemy pawn by moving one square diagonally, left or right, to a square occupied by the enemy. The captured piece is removed from the board. These are the same as pawn moves in chess, except that no double move, en passant capture, or promotion of pawns is permitted.



The game is won in any of three ways:

1. By advancing a pawn to the third row.
2. By capturing all enemy pieces.
3. By achieving a position in which the enemy cannot move.

Players alternate moves, moving one piece at a time. A draw clearly is impossible, but it is not immediately apparent whether the first or second player has the advantage.

To construct HER (Hexapawn Educable Robot) you need 24 empty matchboxes and a supply of colored beads. Small candies that come in different colors—M&M's or Skittles—work nicely. Each matchbox bears one of the diagrams in Figure 35.2. The robot always makes the second move. Patterns marked "2" represent the two positions open to HER on the second move. You have a choice between a center or an end opening, but only the left end is considered because an opening on the right would obviously lead to identical (although mirror-reflected) lines of play. Patterns marked "4" show the eleven positions that can confront HER on the fourth (its second) move. Patterns marked "6" are the eleven positions that can face HER on the sixth (its last) move. (I have included mirror-image patterns in these positions to make the working easier; otherwise 19 boxes would suffice.)

Inside each box place a single bead to match the color of each arrow on the pattern. The robot is now ready for play. Every legal move is represented by an arrow; the robot can therefore make all possible moves and only legal moves. The robot has no strategy. In fact, it is an idiot.

The teaching procedure is as follows. Make your first move. Pick up the matchbox that shows the position on the board. Shake the matchbox, close your eyes, open the drawer, remove one bead. Close the drawer, put down the box, place the bead on top of the box. Open your eyes, note the color of the bead, find the matching arrow and move accordingly. Now it is your turn to move again. Continue this procedure until the game ends. If the robot wins, replace all the beads and play




Figure 35.2. Labels for HER matchboxes. (The four different kinds of arrows represent four different colors.)

again. If it loses, punish it by confiscating only the bead that represents its *last* move. Replace the other beads and play again. If you should find an empty box (this rarely happens), it means the machine has no move that is not fatal and it resigns. In this case confiscate the bead of the preceding move.

Keep a record of wins and losses so you can chart the first 50 games. Figure 35.3 shows the results of a typical 50-game tournament. After 36 games (including 11 defeats for the robot) it has learned to play a perfect game. The system of punishment is designed to minimize the time required to learn a perfect game, but the time varies with the skill of the machine's opponent. The better the opponent, the faster the machine learns.




Figure 35.3. Learning curve for HER's first 50 games (downslant shows a loss; upslant, a win)

The robot can be designed in other ways. For example, if the intent is to maximize the number of games that the machine wins in a tournament of, say, 25 games, it may be best to reward (as well as punish) by adding a bead of the proper color to each box when the machine wins. Bad moves would not be eliminated so rapidly, but it would be less inclined to make the bad moves. An interesting project would be to construct a second robot, HIM (Hexapawn Instructable Matchboxes), designed with a different system of reward and punishment but equally incompetent at the start of a tournament. Both machines would have to be enlarged so they could make either first or second moves. A tournament could then be played between HIM and HER, alternating the first move, to see which machine would win the most games out of 50.

Similar robots are easily built for other games. Stuart C. Hight, director of research studies at the Bell Telephone Laboratories in Whip-