

Activity 7

The Sweet Computer: machines that learn

Age group successfully used with:	9 – 18
Abilities assumed:	none
Time:	approx 50 minutes
Size of group:	3-30 Groups of around 12 per facilitator work well, giving everyone lots of chances to be involved

Focus

How can a computer learn from experience?

Summary

The class challenge a machine made of sweets to the game of Hexapawn (a little like a game of chess with only pawns and on a 3x3 board). At the start the machine only knows the basic rules of the game. The more it plays the better it gets: when it loses the class punish it by eating its sweets allowing it to learn from its mistakes.

Aims

The aim of this activity is to demonstrate how a computer can learn. Intelligence is more than just the skills you have. Intelligent creatures can adapt and learn from their experiences, but so can computers. Even though they follow fixed rules, those rules can allow them to learn from their mistakes and adapt their behaviour.

Technical Terms

Machine learning.

Materials

For each person

- Rules of Hexapawn sheet (provided)
- 1 sheet of squared paper

For each sweet computer

- Large flip chart sized squared paper
- Large noughts and crosses board with 3 “X” pieces and 3 “O” pieces
 - The ideal is the large garden play ones that can be bought from toy shops. Alternatively create one from coloured paper and frisbees.
- Board position cards (provided)
- 24 Plastic cups
- Lots of sweets of different colours matching the arrows on the boards.

Created by Paul Curzon, Peter McOwan and Jonathan Black, Queen Mary University of London based on the Matchbox Computers of Donald Michie and Martin Gardner with support from EPSRC and Google

If using the supplied board sheets you will need for each session to teach the computer from scratch

- 5 red sweets
- 11 yellow sweets
- 13 purple sweets
- 12 brown sweets
- 13 blue sweets
- plus lots of spares of any colour to be eaten after each game

What to do

Preparation: Before the session you will need to buy lots of sweets! Clear a large floor or table space and place the 3x3 playing board at one end with pieces laid out as shown in Figure 1. Leave enough room for people to move round the board.




Figure 1: The initial board position – X (the human) plays first

Now lay out the Sweet computer at the X end of the board. Place the separate board position cards (provided) so that those for Move 1 are together as are those for move two, etc. This will make it easier to find the picture with the current board position quickly.

Place an empty cup next to each board picture. Set up each cup in turn as follows. For each arrow on that board, place one sweet with the same colour as the arrow in the cup. For example, the cup that goes with the board position of Figure 2 should have 1 yellow and 1 red sweet. The cup that goes with the board position of Figure 3 should have 1 blue and 1 purple sweet.

The Sweet computer is now set up to play.

Rules of the Game

The game our sweet computer is going to teach itself to play is called Hexapawn. It's a little like playing a mini-game of chess with pawns only. It is played on a 3x3 board.

Created by Paul Curzon, Peter McOwan and Jonathan Black, Queen Mary University of London based on the Matchbox Computers of Donald Michie and Martin Gardner with support from EPSRC and Google

3 X pieces and 3 O pieces (or black and white) are placed on the first and last rows of the board as shown in Figure 1. To make a move you can either:

1. **Move one place forward if nothing is there, or**
2. **Take an opponent's piece that's in the next place diagonally**

For example, from the position in Figure 2 there are two moves that the left-hand X can make. It can move forward into the space or take the O diagonally. The central X cannot move however. It can't move diagonally as there is nothing to take, and cannot move forward as there isn't a space to move into.

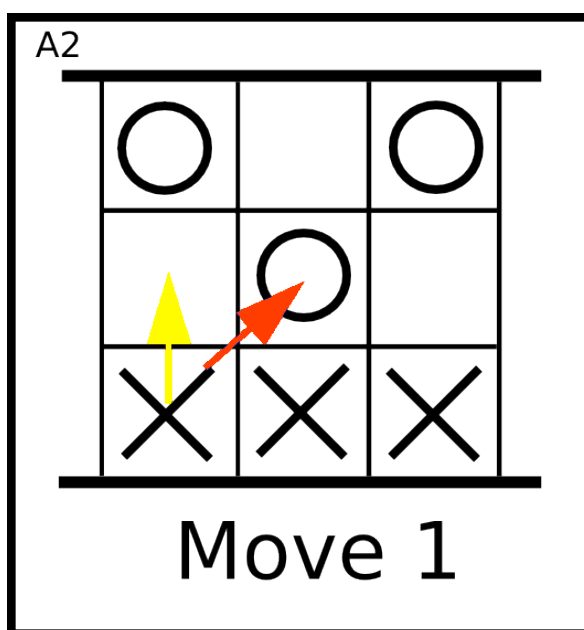


Figure 2: An example board position showing possible moves for X (with symmetries removed)

How do you win? There are 3 ways:

1. **Getting a piece onto the last row,**
2. **Fixing things so the other player can't go at all, or**
3. **Taking all the other player's pieces**

For example, in Figure 3, both moves shown win. If the front X moves it reaches the last row so wins. If the back X moves, it takes the last O so also wins.

That's all there is to it. Play a game as you explain the rules to everyone to get the idea.




Figure 3: Some winning moves

Playing the Sweet computer

Get a group of 3 out at a time to play each game, cycling through the class. One person plays for the humans, one person makes the Sweet computer's move and the third person operates the Sweet computer – working out what move it wants to play.

One cup matches each board position that the Sweet computer could end up in. One sweet matches each legal move it could possibly make and so are used to decide which move it will make. Here's how below.

How the Sweet computer plays

Humans go first, and are the noughts so... the person playing for humans makes their move.

To keep the size of the Sweet computer small, for the first move (only) we have doubled up two possible moves. If the person makes their first move down the RIGHT edge, then swap it for the equivalent move down the LEFT edge. Point out that this will make no difference to the game.

Now it is the Sweet computer's turn. Each time it wishes to make a move do the following:

1. Find the board position card that matches the position the board is actually in after the human's first move.
2. The person operating the Sweet computer should pick up the cup on that board card and close their eyes. They shake the cup and take out a sweet at random.
3. Put the sweet on top of the board position. Note its colour and find the arrow with the same colour on that board. That arrow shows what move the Sweet computer wants to make. Leave the sweet where it is next to the board until the end of the game.

4. The player making the moves for the Sweet computer makes that move.

Now it is the human's turn again, so repeat the above.

If you ever get into a position with no sweet in the cup corresponding to the current position when it is the Sweet computer's move, then the sweet computer resigns – the human wins, but if you can't go or the sweet computer gets a piece to the end then it wins!

Play a game or two with the first set of three with you explaining exactly what each person should do step-by-step. They will quickly get the idea.

If the human gets to a winning position then, with a little help from the class, it will learn from its mistakes.

Learning from its mistakes

Here is how the Sweet computer learns. Remember at the end of a game there will be a sweet next to several board positions showing the moves made.

1. If the computer lost then the human who beat it eats the sweet that is on the last board position that made it play its final losing move. It will never play that losing move again!
2. If that cup is now empty then put it aside. The Sweet computer will resign if it gets to that position again.
3. If the Sweet computer won then you don't eat that sweet. (With younger children it is a good idea to have a supply of spare sweets so that everyone gets a sweet each time whether they win or lose – it turns out that children do not respond to Sweet based punishment as well as the Sweet computer!).
4. Put all the remaining sweets back in the cup they came out of.

Now play games over and over again. Swap the three people after each game for a new set of three, rotating through the class. You will probably need to play around 20 games for it to start winning consistently, though this depends on how the humans play. Get the class who are not playing to try and identify the sequence of moves to play that could still beat it while they are waiting for their turn. This works well in groups of around a dozen as then everyone is likely to get to play all the parts.

Eventually the Sweet computer should narrow down on one move as the correct first move, with no sweets left in the other alternative. When this happens point out that it has now decided where to go for that move.

Remember. If it loses, punish it - eat its sweets!

Charting its learning progress

An easy way to get the class to record the progress of the Sweet Computer as it learns is to create a graph (see Figure 4). Use squared paper – give everyone a sheet to record progress and also make one person the official recorder with a giant sheet of flipchart squared paper. Before the first game get everyone to draw a horizontal (zero) line across the middle. When the Sweet computer loses a game everyone draws a

diagonal downwards across one square. Similarly if it wins they draw a diagonal sloping upwards.

Figure 4 gives a picture of the way the computer learns. Initially the lines are likely to be heading down. As it starts to learn, it will win more and more and the line will head upwards. When below the zero line the Sweet computer is losing more than it is winning and when above it is winning more than losing. The shape also gives an idea of how fast it learns.

At the end compare the charts and discuss what they show.




Figure 4. Record of wins (line up) and losses (line down)

What's going on?

The sweet computer will play really badly to start off. It's just making moves at random. The more it plays and the more it loses so the more sweets you eat, the better it will get. Eventually it will play perfectly. No one told it how – it learnt from its mistakes – and because the class ate its sweets!

It's possible to write a computer program that learns in the same way – but without sweets. Instead, the sweets are just 1s in the computer's memory that are flipped to 0s. Not so yummy!

Variations and Extensions

This activity works well in an hour session using the “intelligent piece of paper” activity as the introduction. That is likely to lead to the suggestion that the paper isn't intelligent as it doesn't learn. This activity then demonstrates how that can be done too.

Reward and Punishment

The class can do some experiments to investigate what the best way to teach the sweet computer is. Do this by trying slightly different rules.

At the moment we just punish bad moves (by eating the sweets). What about when the sweet computer wins? You could reward all the moves that led to it by not only putting the sweet back in the cup but by adding another of the same colour too. Now the Sweet computer will be more likely to make those moves again.

Get the class to try out other variations of rewards and punishments.

Decision Trees and Rules

If you have played enough games so that the Sweet computer becomes unbeatable then as a harder exercise the class could be set the task of analysing its memory to try and write down a set of rules that could be followed to play perfectly. What should a perfect player do for the first move? What should they do for the second? etc. See the “intelligent piece of Paper” activity as an example of such a set of rules for playing perfect noughts and crosses.

One way to start to do this is to remove the board diagrams of all positions with no sweets. This gives just those board positions you need to consider. Put them into a decision tree structure ie draw a tree diagram with nodes board positions and edges the move the human might play.

Next write rules to distinguish which moves should be made. The rules will probably be of the form:

IF *the board position is like this* THEN *make this move*

Once the group have a set of rules, get someone to follow them and see if they are both clear enough that the person makes the correct move and correct in that they do win.

Further Reading

A Boost for learning.

How machine learning can be used to beat SPAM

<http://www.cs4fn.org/optimization/boost.php>

The Machines can translate now

Machines that learn to translate

<http://www.cs4fn.org/linguistics/statstranslation.php>

Go, Go gadget: computers' biggest gaming challenge

Progress at programming the last game humans are better at

<http://www.cs4fn.org/ai/gogogadget.php>

Links to other activities

The intelligent piece of paper

Take part in a test of intelligence against an intelligent piece of paper!

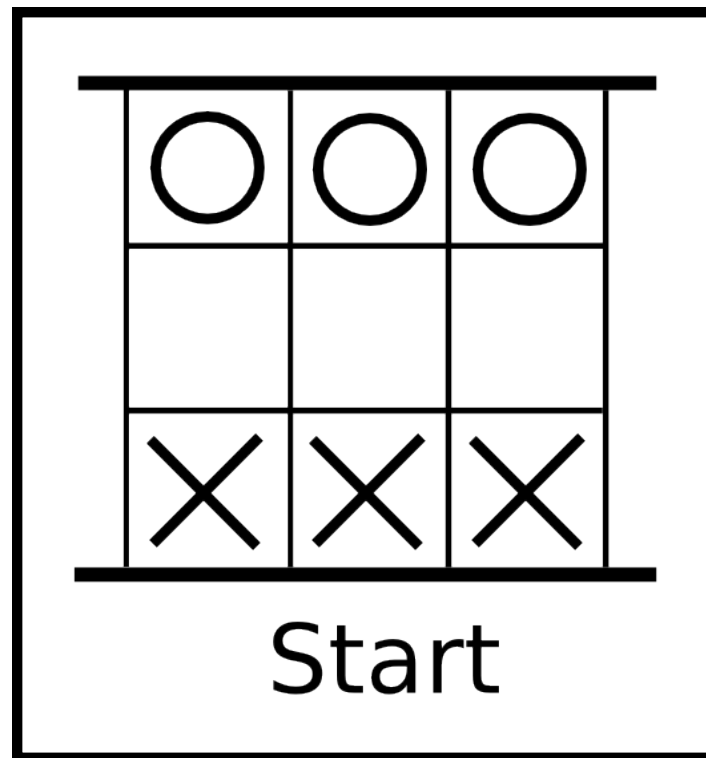
This is a good introduction to what a computer program is, and also to start a discussion on what it would mean for a computer to be intelligent.

Winning Games: the perfect Tic-tac-toe player

Create a set of instructions that would allow anyone to play Tic-tac-toe perfectly.

This is a good follow on activity from the intelligent piece of paper – now create your own! It introduces programming and explores how a computer is able to win at board games like chess. The emphasis is on programming being about solving the problem rather than about being able to write in a programming language.

The Rules of Hexapawn



The game of Hexapawn is played on a 3x3 board. 3 black pieces and 3 white pieces (or Xs and Os) are placed on the first and last rows of the board as shown. To make a move you can either:

3. Move one place forward if nothing is there, or
4. Take an opponent's piece that's in the next place diagonally

It's a little like playing a mini-game of chess with pawns only.

How do you win? There are 3 ways:

4. Getting a piece onto the last row,
5. Fixing things so the other player can't go, or
6. Taking all the other player's pieces

That's all there is to it. You might want to play a game to get the idea.