CS4FN

# artificial intelligence!

## ...but where is the intelligence?

**Paul Curzon, Peter McOwan and Jonathan Black**

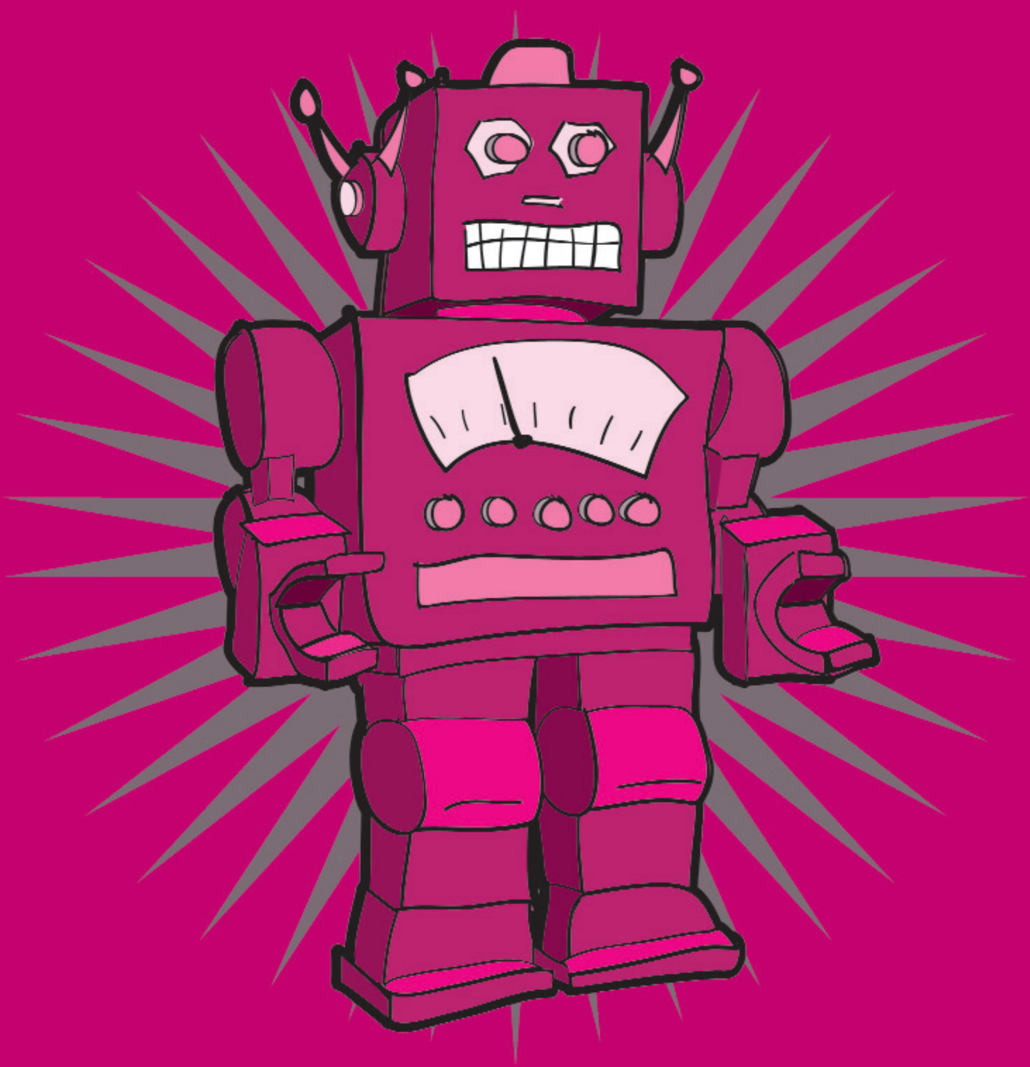

# Contents

# introduction

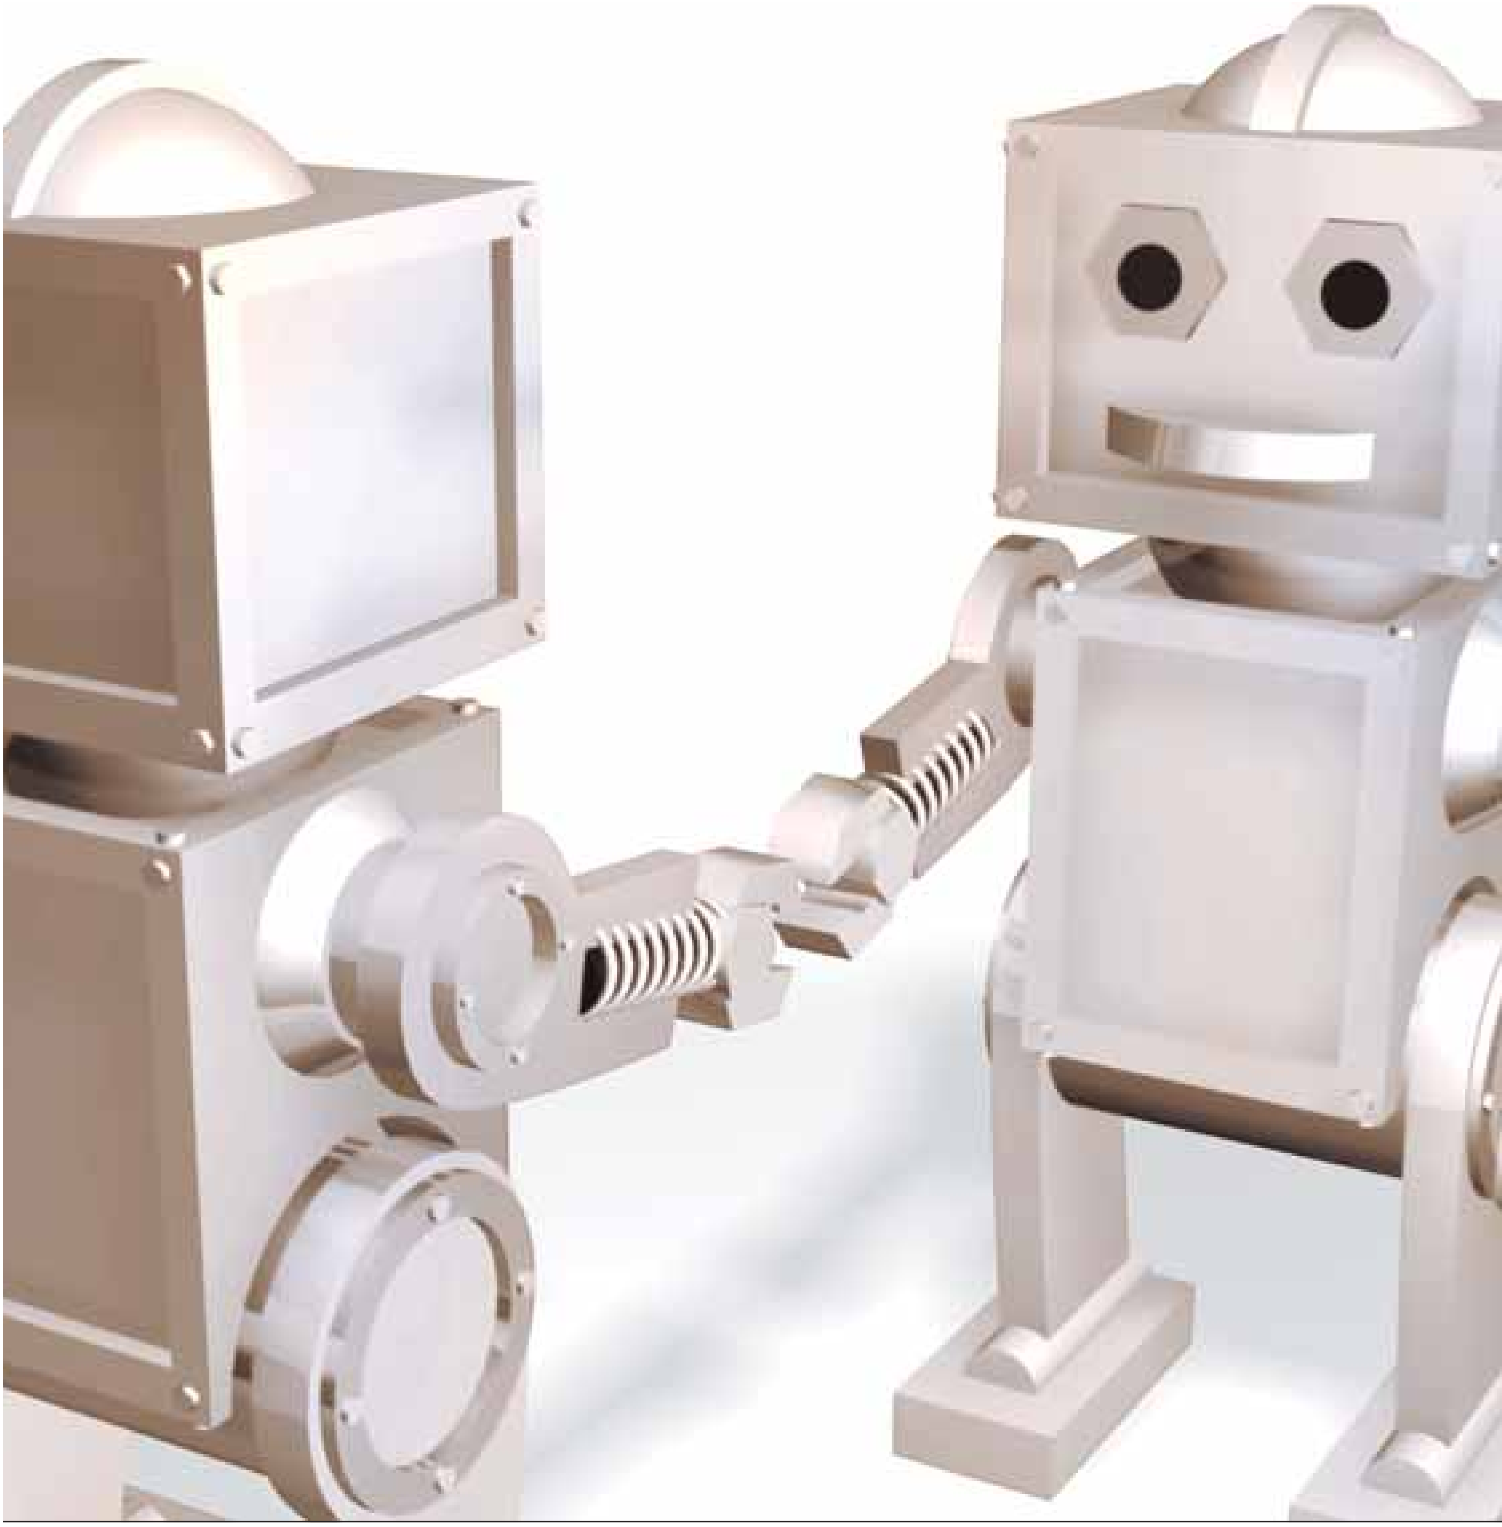

# artificial intelligence?

Artificial intelligence – machines that think. Is it really possible or is intelligence something only humans have? How would you spot an intelligent computer? How would you build one? What is intelligence anyway?

# living dolls in history

## Through history inventors have strived to make 'living dolls' – machines that could act like real creatures and do intelligent things – and back then they had to do without computer electronics.

**Greek pigeons**
350 BC: a Greek mathematician called Archytas of Tarentum built a model mechanical bird called 'the pigeon'. It was pushed along a string by steam.

**Italian knights in armour**
1495: Artist and inventor Leonardo da Vinci designed a mechanical robot that looked like a knight in armour. The cogs and gears inside made it move as if someone was actually inside it.

Lots of inventors in medieval times designed mechanical creatures to amuse their kings and queens.

**French musicians**
1738: French inventor Jacques de Vaucanson built a series of mechanical creatures called automata. His first was a model of a flute player. It could play twelve songs on a real flute. His second attempt was an automaton that played a flute and a drum or tambourine.

**Swiss dolls**
1770: Two Swiss clock makers, Pierre Jaquet-Droz and his son Louis, built three mechanical dolls. One could write, one could play music and the third could draw. These lifelike dolls entertained royal families all over Europe.

As if that wasn't enough Pierre also found time to invent the wristwatch.

## duck!

Jacques de Vaucanson also made a duck automaton that could move, quack, flapped its wings and even ate and digested food! (It pooed too!)

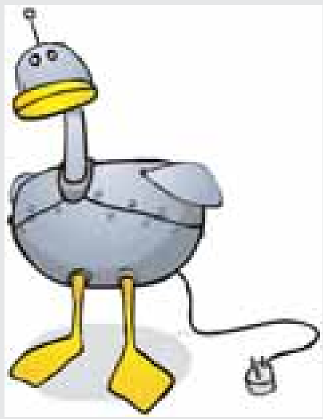


how do
you
recognise
intelligence?

# What is intelligence?

## What can we do that shows we're intelligent?

There are lots of abilities we have that show our intelligence. You might be good at reading but your friends are good at puzzles. She has an amazing memory for facts and he is good at understanding people. We are all intelligent but in different ways.

**What can people do that 'shows' they are intelligent in some way or other?**
It's supposed to be our intelligence that makes we humans stand out, but animals can be clever too. There are lots of things that people have said no animal can do, but then someone finds an animal that can do it!

**How many things can you think of that show us that animals are clever too?**
Start by thinking about pets – cats, dogs and mice. Compare them to a doll – can they do clever things that help them survive without needing to be looked after?

Now think of other animals like chimps or crows, bats or elephants. Have you heard of any amazing abilities they have that show they are intelligent?

Perhaps there are things they can do that people can't do at all!

# almost human?

Is it possible to put animals in order of intelligence? It isn't as easy as you might think. Different animals (including humans) are intelligent in different ways, just as different people are.

- Think about how intelligent the animals (and other things) on the next page are.
- Which would you say are intelligent at all?
- Which do you think have the most human-like intelligence?
- Which do you think have the least human-like intelligence?
- Do any have intelligence in ways humans don't?

## A mouse
I can find food and avoid being eaten by others. I can learn the way round mazes to find food at the centre.

## A pet rock
I'm just a stone with eyes and mouth painted on. I can't do anything much at all, but am easy to get on with.

## A dog
I can find things to eat with my great sense of smell. I love to play and will protect my owner. I can learn tricks and even act as a guide dog for a blind person, helping them cross roads.

## 20Q
I can read your mind! Well, almost. I can quickly work out what you are thinking of just by asking you yes/no questions.

## A human baby
I can't walk, talk, find food or even see very well, but I will soon learn to do all these things.

## A robot vacuum cleaner
I can move around a room working round objects. I make sure I clean the whole room. When I run out of energy I can get back to my base to 'feed' myself, even if my base has been moved.

## A bee
I can find flowers with pollen, then return to the hive and do a special dance to tell the other bees how to find them too. My dance tells them which direction to go in and how far to fly to get to the pollen.

## An electronic sudoku solver
I can solve logical puzzles – even the hard ones – faster than any human.

## A turtle
After years growing up somewhere else I can easily find the beach I left when I was born to lay my own eggs there.

## A parrot
I can repeat things humans say. I can fly and find food to eat. I know all the other birds in my flock and can recognise a stranger.

## A fly-by-wire jet
I can take off in New York and land at Heathrow airport without a human doing anything.

## A FurReal Squawkers McCaw
I can say lots of things and learn new words by copying what you say. I can dance, sing, whistle and laugh. I react when you pat me or ruffle my feathers.

## A chess-playing computer
I can beat humans at chess, working out good moves and bad moves.

## A cuttlefish
I am a master of camouflage. I can blend into any background, even changing the pattern of my skin to fit. I can also recognise different threats and give different warning signs depending on what the danger is.

## A chimp
I can solve problems using tools – like fashioning a stick to use to 'fish' for ants out of tree trunks. I also share my food with others.

## CCTV with computer vision system
I can tell men and women apart from their faces and the way they walk. I can follow one person as they pass from camera to camera. I can tell whether people are happy or sad from their expressions.

## An ant
I am able to find food then return straight to the nest. I leave a trail that other ants can follow straight to the food. We ants are good managers, dividing up work so that different ants in the colony do different things.

## A homing pigeon
I can find my way home even if I'm taken hundreds of miles away in a lorry. I am naturally good at knowing the right way to go but I also know some extra tricks like following motorways.

# How clever is that?

# iq and all that

## The most common way to test people for intelligence is an IQ test. The trouble with these tests is that they only measure one or two kinds of intelligence.

Worse still, quite often the questions have several answers that could be right and you have to get the one the person who set it thought of! Not such a good test. What different kinds of intelligence are there? Below are some of the ways people can be intelligent.

Circle the three you're best at.

**Logical intelligence**
You can think logically and are good at puzzles like Sudoku and games like chess.

**Musical intelligence**
You are good at music, sounds and beat – you find it easy to learn a musical instrument or sing in tune.

**Verbal intelligence**
You find reading and writing really easy. You may be good at word searches and anagrams.

**Bodily intelligence**
You are good with your hands or feet – you are good at making things and doing crafts or are perhaps naturally good at sport.

**Sensory intelligence**
You are good at seeing things and recognising what they are, or are good at recognising smells or sounds.

**Spatial intelligence**
If you are intelligent in this way then you are good at doing things with pictures like jigsaw puzzles and do not easily get lost.

**Motivational intelligence**
You don't need others to tell you what to do. You know what you want to do and go ahead and make it happen

**Social intelligence**
You understand people really well and are good at helping them understand.

**What's missing from this list? Are there other kinds of intelligence too: perhaps kinds you are really good at?**

**Different artificial intelligences can be intelligent in different ways too.**

**David Beckham can curl a football to land exactly where a striker is going to need it to be to score. He can do it even though the striker is somewhere else when he kicks the ball!**

**What kinds of intelligence does he need to do that?**

# the turing test

## How do you decide if something deserves to be called intelligent? Does it have to pass exams to prove it?

What kind of question would you put in an exam to test to see if a computer was intelligent?

Alan Turing was a mathematician and early computer scientist who came up with one way to decide if a computer was intelligent. It became known as the Turing Test and it's still used today, more than 50 years after he invented it.

### How the Turing Test works

Put a computer and a person in one room, while you go into another. Now ask them both questions by typing them messages. Their answers come up on your screen, and you have to work out which is the human. If after lots of questioning you can't tell which is which, then according to the test you should admit the computer is as intelligent as the human!

There is a million dollar prize for the first computer chatbot to pass a Turing Test. The competition is held every year. So far none have come close to passing it.

Could any animal pass a Turing Test? If not, does that mean no animal is 'intelligent'?

### Party games

Alan Turing got the idea for the Turing Test from a Victorian party game. A man and a woman went into another room and everyone had to work out who was the man. All they could do was write questions on a piece of paper that a referee would take to them and bring back their answers. The woman had to tell the truth. The man could lie and pretend to be the woman.

The Victorians didn't have Xboxes, PlayStations or Wii so they had to make up their own amusements!

The earliest chatbot was a program called ELIZA developed in 1966. ELIZA took the role of a therapist having a conversation with you the patient.

Some people were fooled!

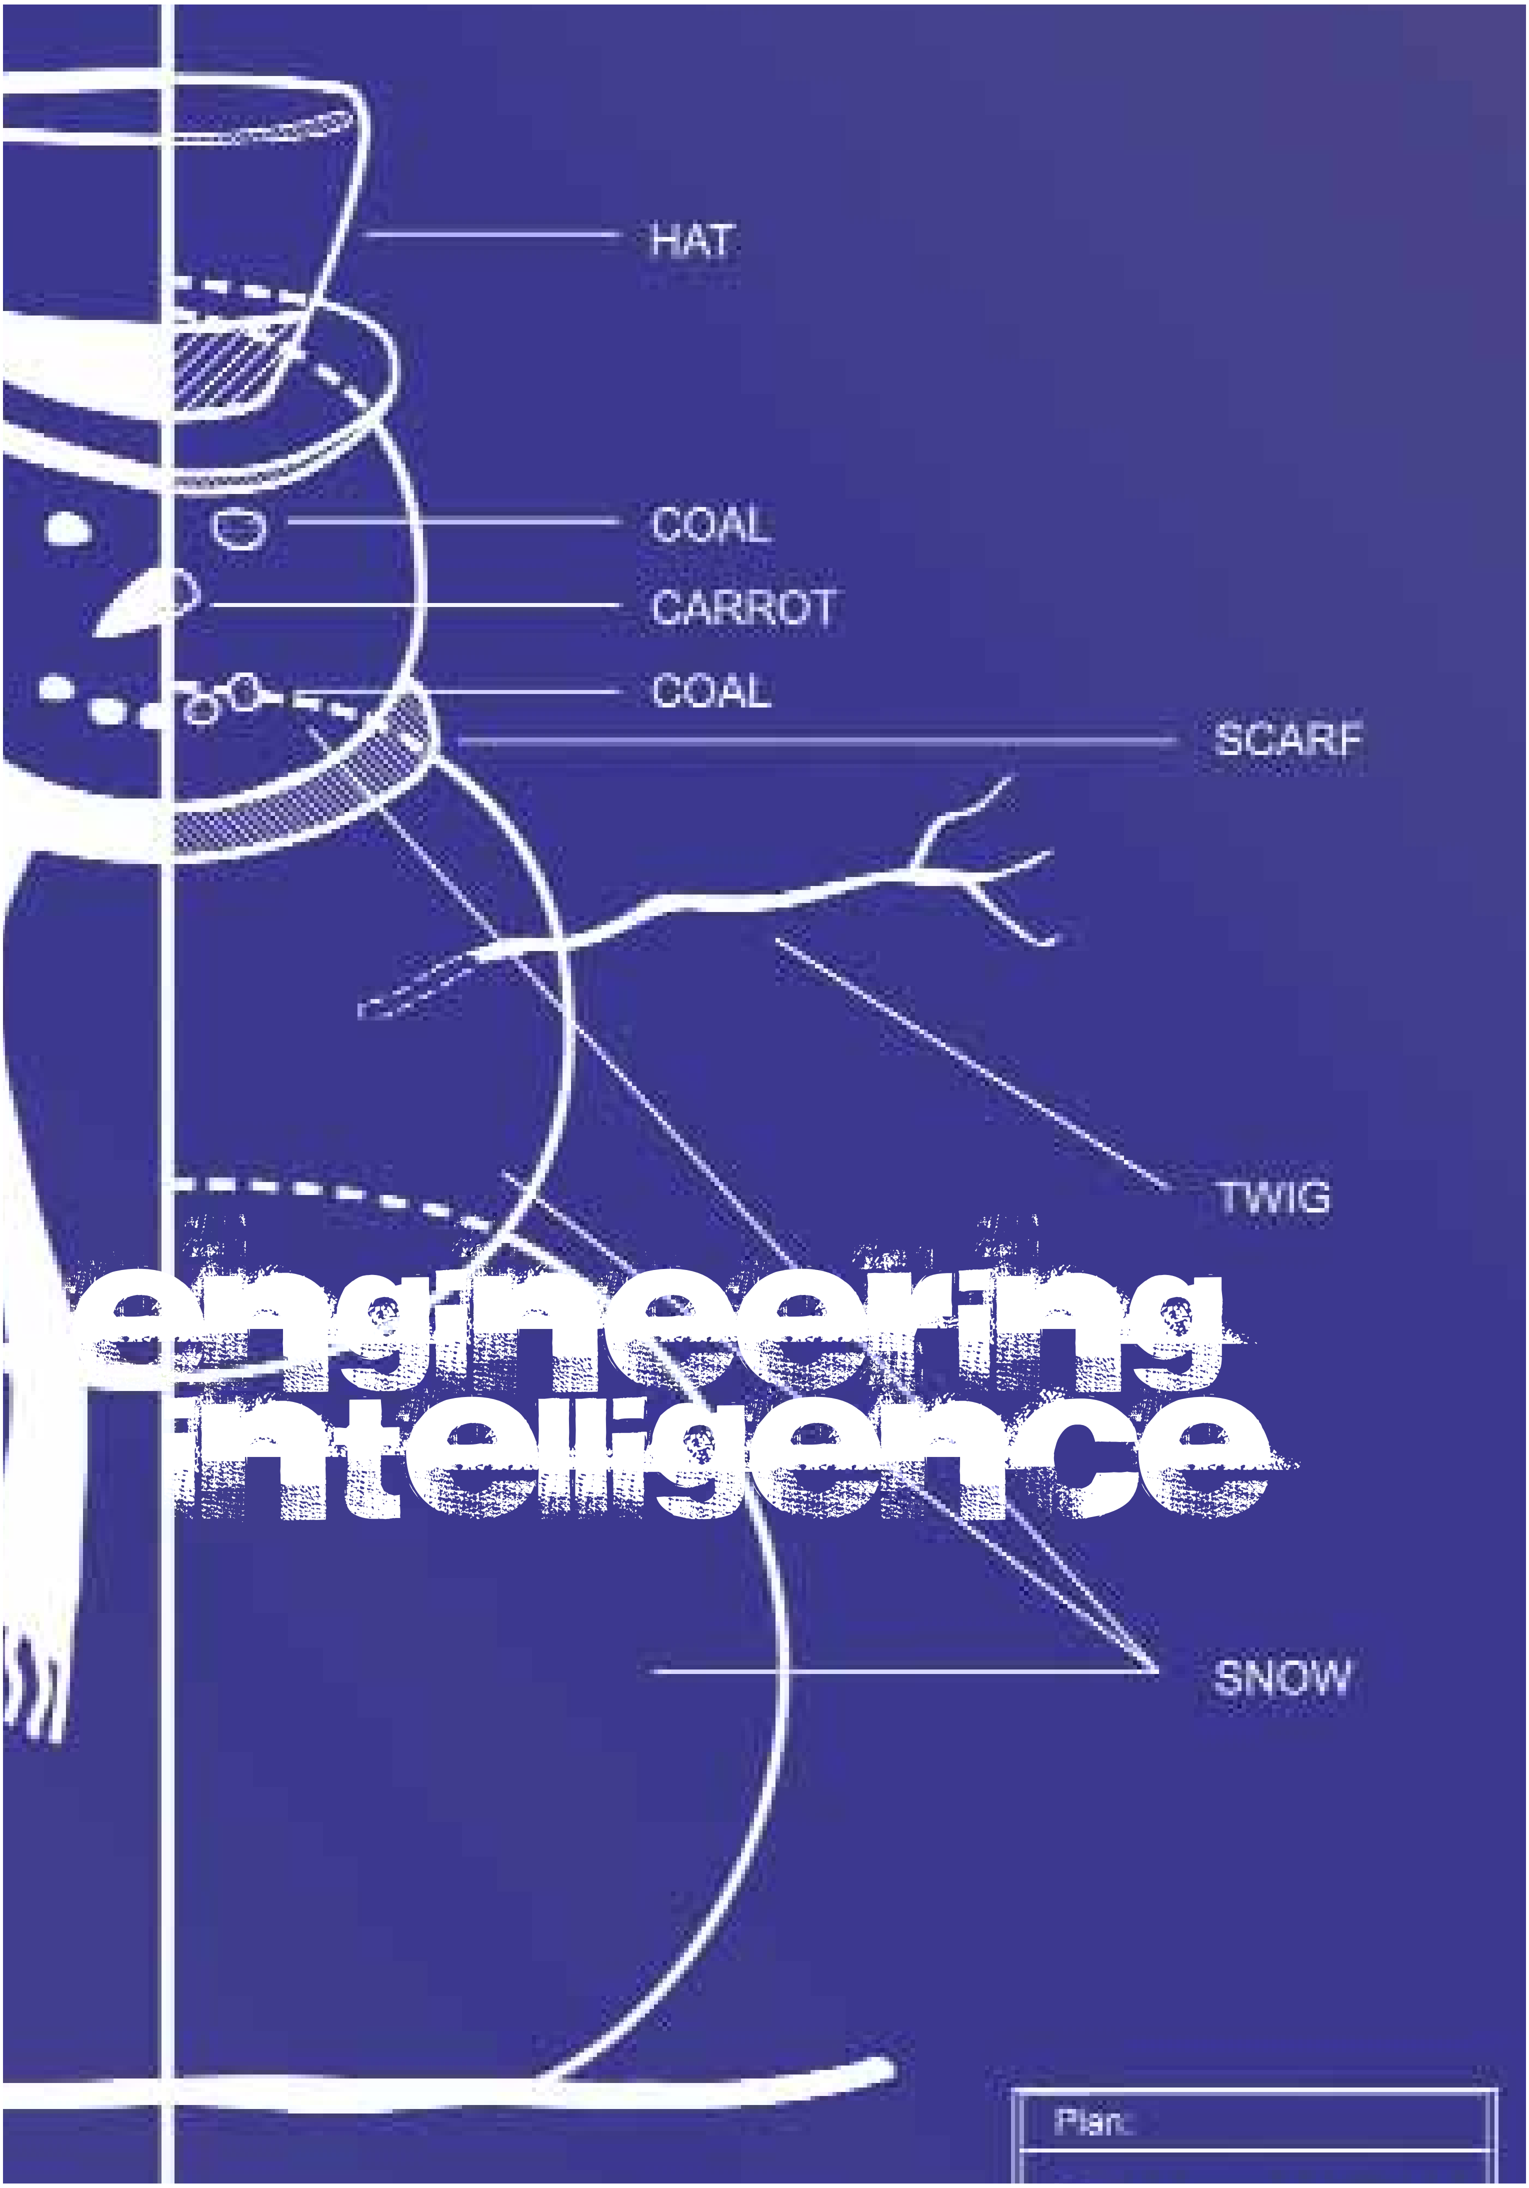
HAT
COAL
CARROT
COAL
SCARF
TWIG
engineering
intelligence
SNOW
Plan:

# an intelligent piece of paper

## This piece of paper is incredibly intelligent.

## Do you believe that? What do you mean you don't?

What reason might there be to claim that a piece of paper is intelligent?

It can't be just that it has some intelligent facts on it – just storing facts is not enough for intelligence. What might this piece of paper be able to do to prove how intelligent it is?

This piece of paper is clever in one way – it is really good at playing noughts and crosses. It has played humans a lot and in fact it has never lost a game. It wins a lot of the time too – and that's despite the fact that if both players play well the game should end in a draw.

Turn over to find the intelligent piece of paper's instructions.

Give it a game, and see how you do. The paper needs you to make its moves for it, but if you follow the paper's instructions exactly you'll see how well it plays noughts and crosses.

The intelligent piece of paper wants to play first. Perhaps you could write some instructions that would allow it to play well if it went second too. Why not have a go?

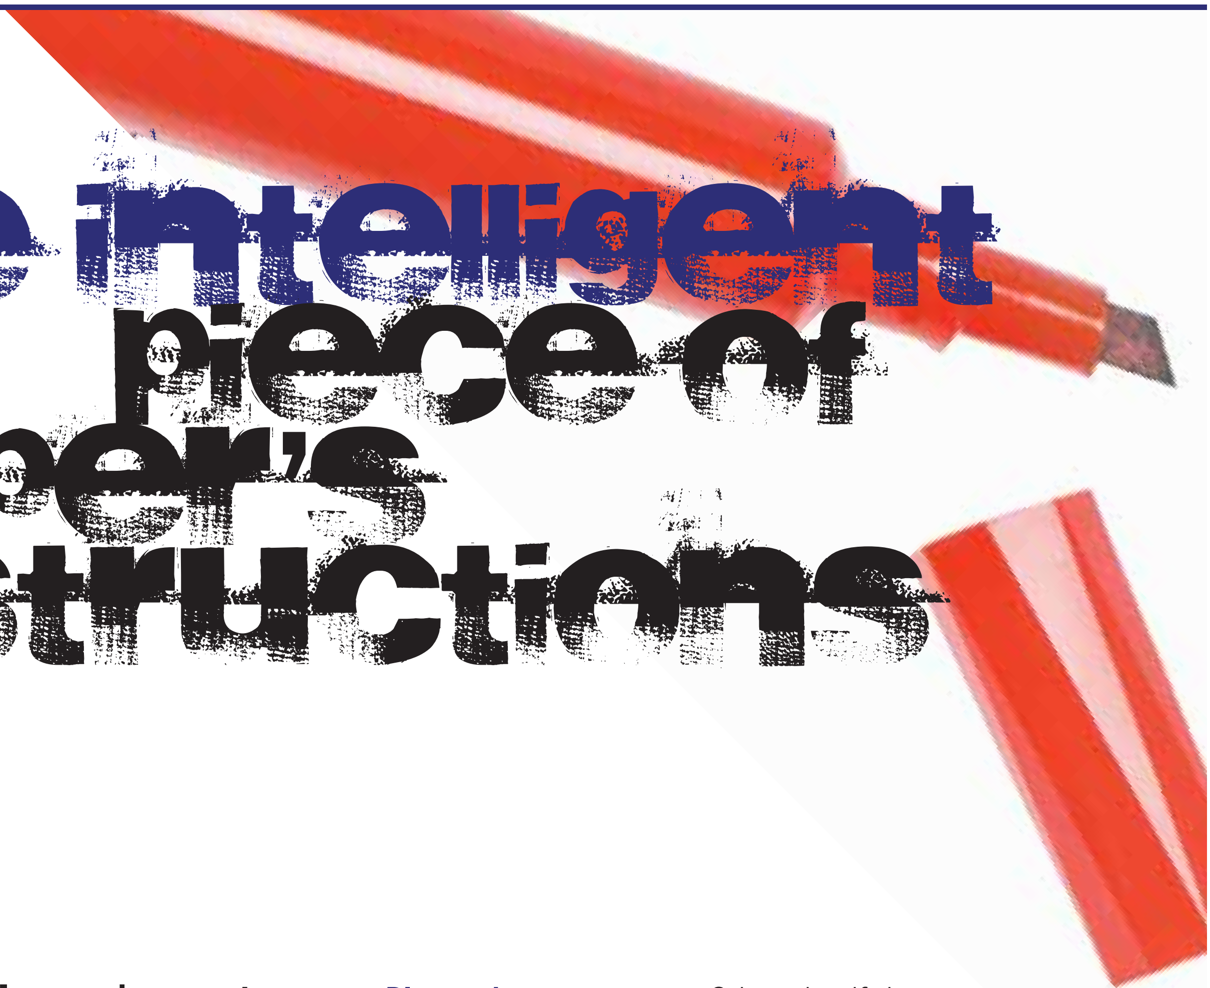

# the intelligent piece of paper's instructions

Hello. It's nice to be playing with you, especially since I'm an incredibly intelligent piece of paper and will probably beat you.

**Please draw me a noughts and crosses grid**
I am X and I go first. Please do the following for me.

**For my first move:**
Draw an X in a corner for me.

**For my second move:**
If you did not go there already then draw an X in the opposite corner to my first move for me.

Otherwise put an X in a free corner for me.

**For my third move:**
If there are two Xs and a space in a line (in any order) then put an X in that space and I win!

Otherwise if there are two Os and a space in a line then put an X in that space. Ha!

Otherwise put an X in a free corner for me.

**For my fourth move**
If there are two Xs and a space in a line (in any order) then put an X in that space and I win!

Otherwise if there are two Os and a space in a line then put an X in that space for me.

Otherwise put an X in a free corner.

**For my fifth move:**
Put an X in the free space for me.

# is rule following intelligent?

Our piece of paper's claim to intelligence is based on what it can do, but all it is doing is blindly following rules. That is all computers do. They follow instructions a bit like the ones for playing noughts and crosses. Everything you have seen any computer do was done like that.

Most people would say the paper isn't intelligent. It is the person, the programmer, who wrote the rules who has the intelligence. That may be true, but even so computers can do lots of clever things just by following rules like that.

Often they can even do things better than the programmers who wrote their rules, like chess computers for example.

Back in the 1950s computers had existed long enough that people started to suggest that one day computers would beat humans at chess. "This is of course nonsense," wrote chess expert Edward Lasker. Perhaps he was sure it took intelligence to play chess well and surely machines just following rules could not out-think humans. But he was wrong…

# Winning at chess

In 1997, computer company IBM's Deep Blue supercomputer beat the reigning world chess champion, Garry Kasparov in a chess tournament. Deep Blue was just a machine following a set of rules like our noughts and crosses instructions. It just takes more instructions to play chess well.

## Who is the smartest of them all?

Presumably because Garry had seen too many Terminator films (and because he thought he was the most intelligent creature on the planet until he lost), he said afterwards:

"I'm not afraid to admit that I'm afraid."

## The mighty fall

Kasparov went on to become a Russian politician. Some say Deep Blue is doing airline reservations in an American airport. So the next time the computer double books you on a flight or loses your bags, remember it just could have been the best chess player on the planet who did it!

…and maybe if the computers get too smart for their own good, we could just punish them by giving them dull jobs to do. Ha!

## The computer revolution

Perhaps giving Artificial Intelligences (or AIs) dumb jobs to do is not such a good idea. It's that sort of treatment that has led in the past to workers overthrowing their rulers – so maybe that's how the real computer revolution will start!

## Beating your maker

Kasparov was the best chess-playing human at the time. He could easily beat any of the programmers who wrote Deep Blue's instructions. The computer could play far better than its programmers!

## Go! Go! Go!

The Chinese game of Go is actually much harder than chess. So far humans are much better than computers at playing Go.

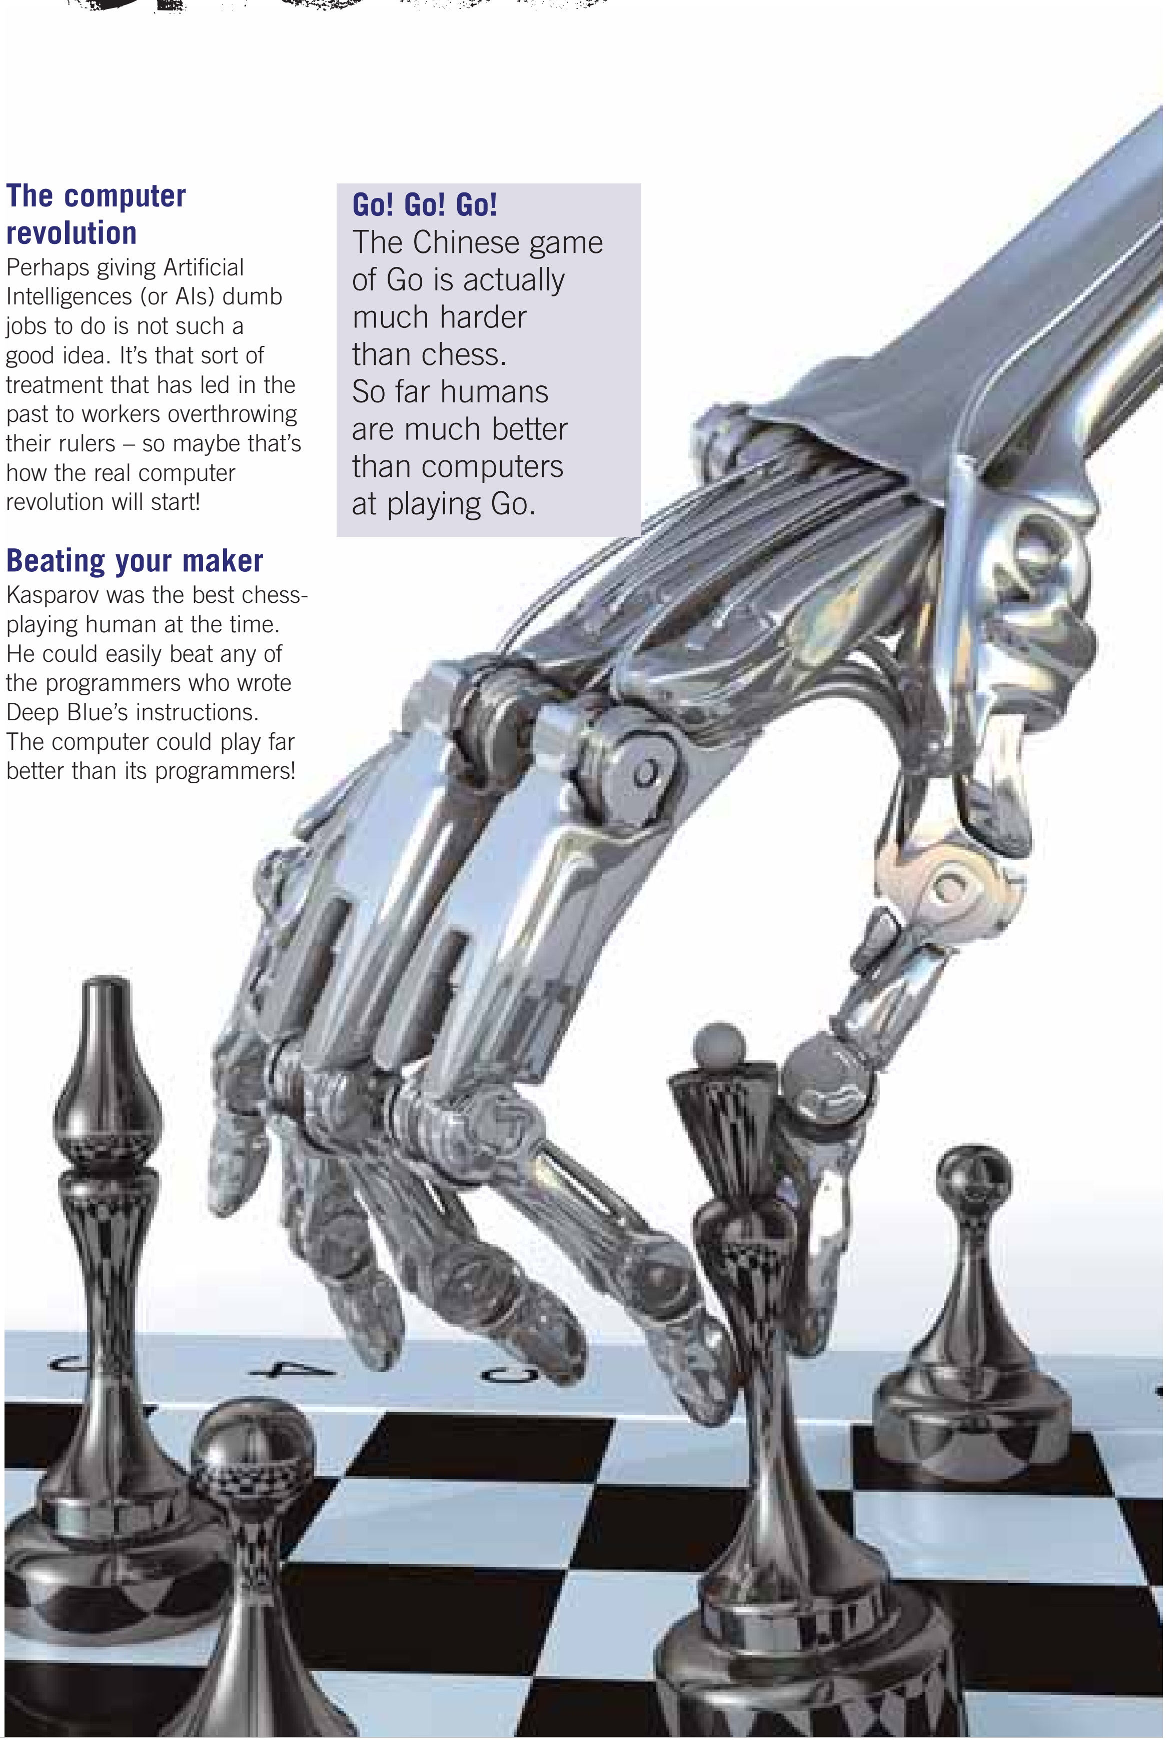

# experts and patterns

A chess computer plays the game in a different way to a human player. It plays lots of pretend games before making each move, trying out as many moves as possible and following the game through, trying different moves for the other player too. The computer then picks the move that gives it the best chance of winning.

Good chess players like Kasparov will look a few moves ahead, but nowhere near as far as a chess computer. It is too hard to do and would take too long. Human players instead look for patterns in the game and know from past experience the best kind of move to make given the pattern.

Experts in other areas, like firefighters, doctors or musicians, do the same thing. Experts don't think through all the hundreds of possible choices they could make in any situation before deciding which is best. Instead they recognise the situation and come up with a solution. When they think through how it will work, if it seems good enough, then they go with it. Their expertise comes from having seen so many situations in the past that their first solution, based on their experience, is usually very good.

Some scientists think that's what 'intuition' is – instantly matching the current situation to something similar the expert has seen before.

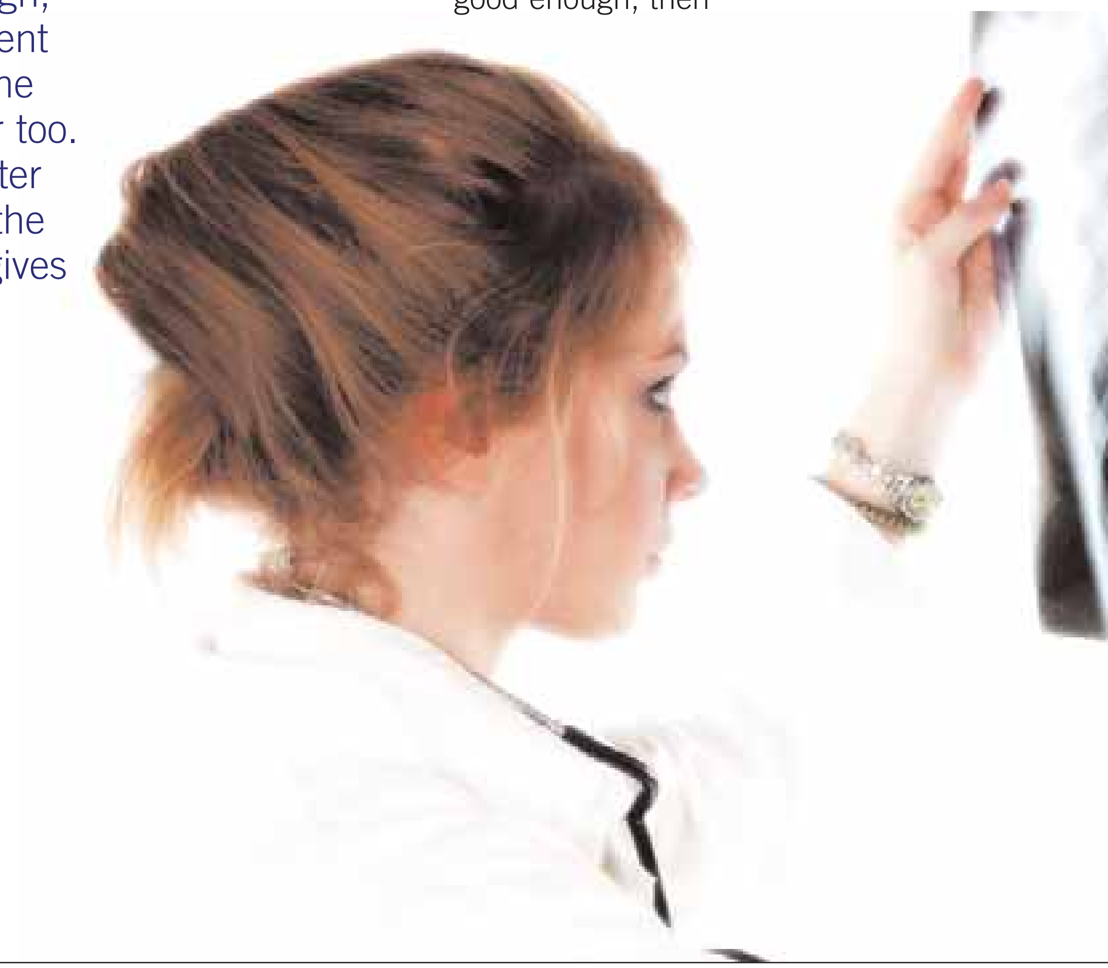

# are you adaptable?

## Our noughts and crosses-playing computer was really good at playing the game, but only when it got to go first. What happens when it goes second? Play against it for a few games and see.

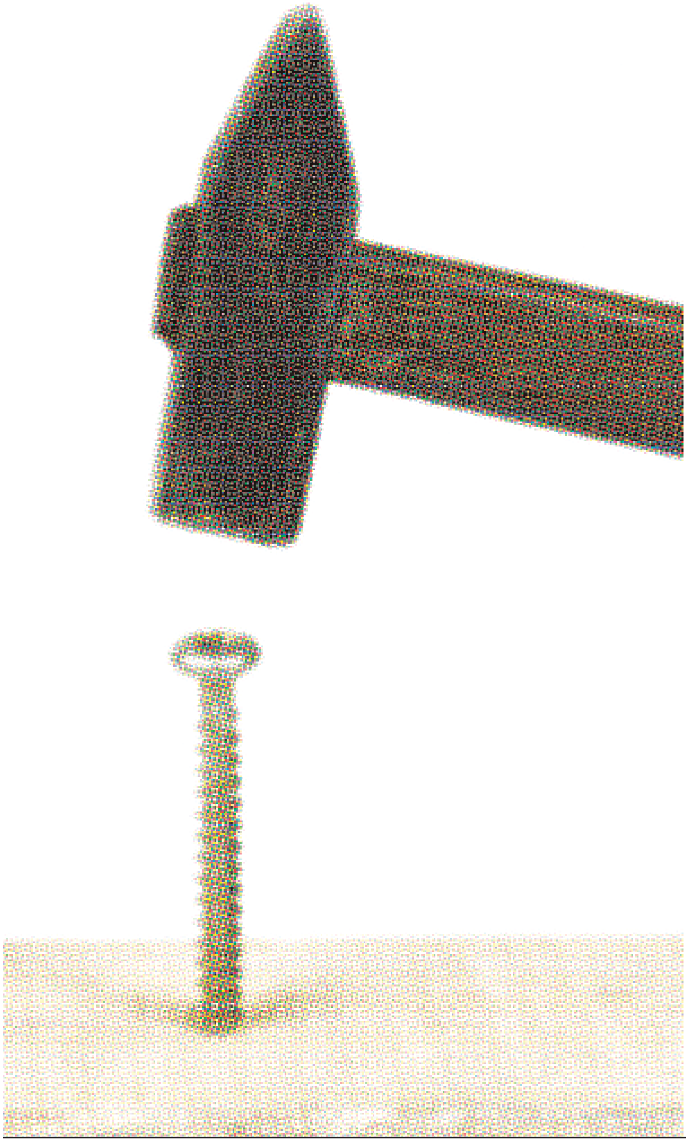

It probably looked as though it was pretty stupid. The computer's suddenly very easy to beat! Its rules were written for one situation and work great in that situation only. Change things just a little bit and suddenly things are not so good.

Humans show their intelligence by being adaptable. Once you've learnt to play well going first you would probably play OK second too.

Fixed rules like our noughts and crosses instructions aren't adaptable. They do what they do and nothing more. Artificial Intelligence researchers call this the 'Frame Problem' – the solution only works for the problem it was made for. It's easy to think of a list of things a computer must be able to do to be intelligent. People can then write programs to do each thing – play chess, recognise faces and so on. The trouble is, as soon as the computer has to deal with a problem the programmers didn't think of, the computer gets lost.

To show really intelligent behaviour, the rules written must somehow be able to cope with new situations. They must also be able to learn from their experiences.

# learning to win

Intelligent creatures are adaptable. They can learn from their mistakes. Could a computer do that? Could it learn just by following rules? Well actually, yes.

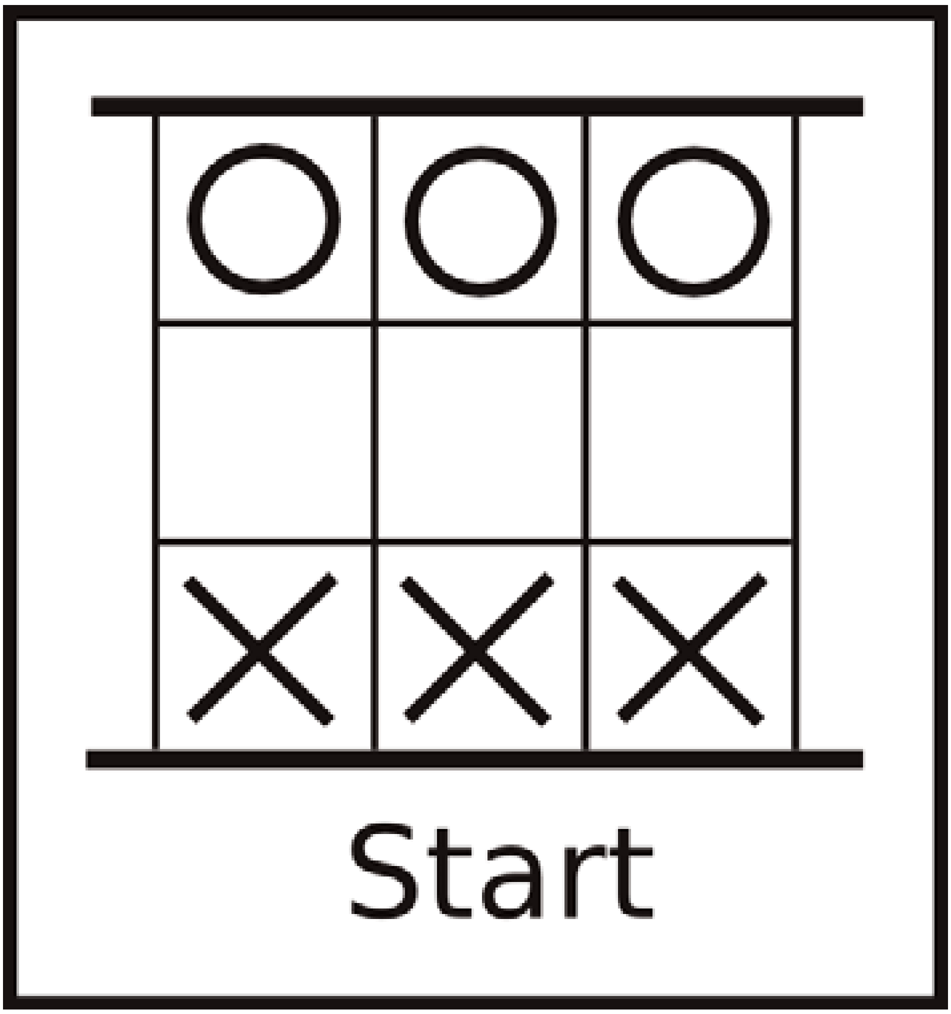


Let's play a game called Hexapawn to show you how. Here are the rules.

The game is played on a 3x3 board. 3 'x' pieces and 3 'o' pieces are placed on the first and last rows of the board as shown. 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

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

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, or
3. 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.

# a sweet computer

## It turns out you can make a learning 'computer' out of boxes and sweets. It starts off stupid but after playing Hexapawn for a while, gradually learns to play better and better...until it never loses. Here's how:

**Ingredients**

You need:

- 24 boxes
- large numbers of coloured sweets of at least 4 colours
- a copy of the 24 board positions shown on pg 22
- A blank 3x3 board to play on
- 3 'x' pieces and 3 'o' pieces to play the game with
- Coloured pens or pencils

**Building your sweet computer**

1. Copy out the 24 pictures of the board shown and stick one on each box.
2. Colour in the arrows so that each is a different colour (the colours you use should match the colours of sweets you have).
3. Take a box and put one sweet in that box to match each arrow on its board picture: if there are three arrows coloured red, green and orange on the box put a red sweet, a green sweet and an orange sweet in the box.
4. Fill all the boxes in this way to match the arrows on the box's board position.

Now your computer is ready to play.

### Play the game

One box matches each board position that our sweet computer could end up in. The sweets are used to decide which move it will make. Here's how:

### How the sweet computer plays

You go first, and are the noughts so…

1. Play a move.
2. Find the box that shows the position the board is in.
3. Close your eyes. Shake the box and take out a sweet at random.
4. Put the sweet on top of the box. Note its colour.
5. The sweet computer makes the move shown by the arrow with the same colour as the sweet.
6. Go back to step 1. It's your turn again.

If you ever get into a position with no box then the sweet computer resigns – you win, but if you can't go or the sweet computer gets a piece to the end then you lose!

### At the end of a game

With a little help from you after each game the sweet computer learns from its mistakes. Here is how. Remember at the end of a game there should be a sweet on top of several boxes.

1. If the computer lost then eat the sweet that is on the last box that made it play the last losing move. It will never play that losing move again!
2. If that box is now empty then put it aside. The computer will resign if it gets to that position again.
3. If the computer won then you don't eat a sweet.
4. Put all the remaining sweets back in the box they came out of.
5. Now play over and over again.

‘If it loses, punish it - eat its sweet!’

## copy, cut out and play!

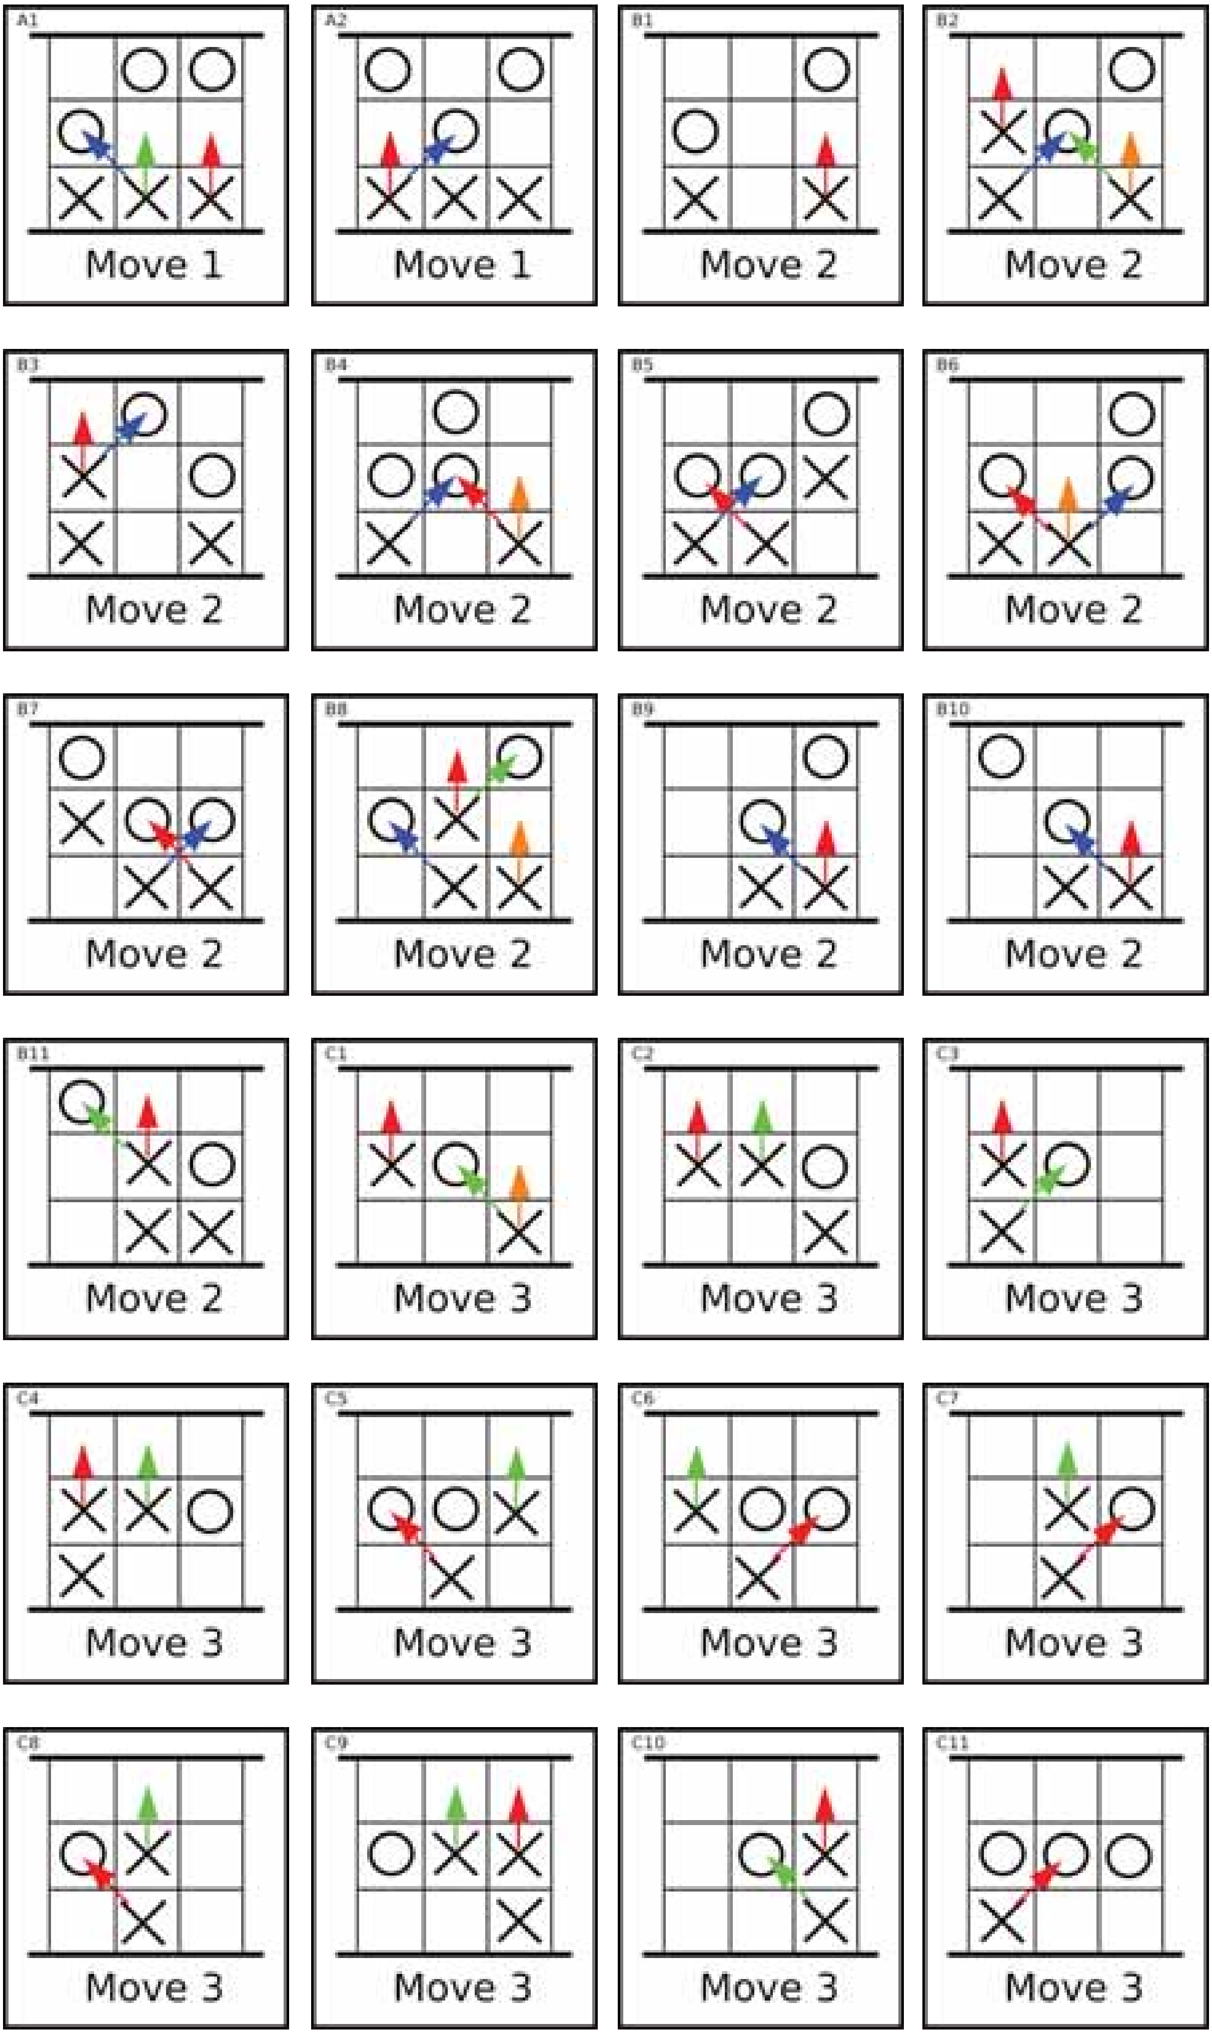

## What's going on?

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 you 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!

You can do some experiments. Try 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 box but by adding another of the same colour too. Now the sweet computer will be more likely to make those moves again.

What other variations of rewards and punishments could you try?

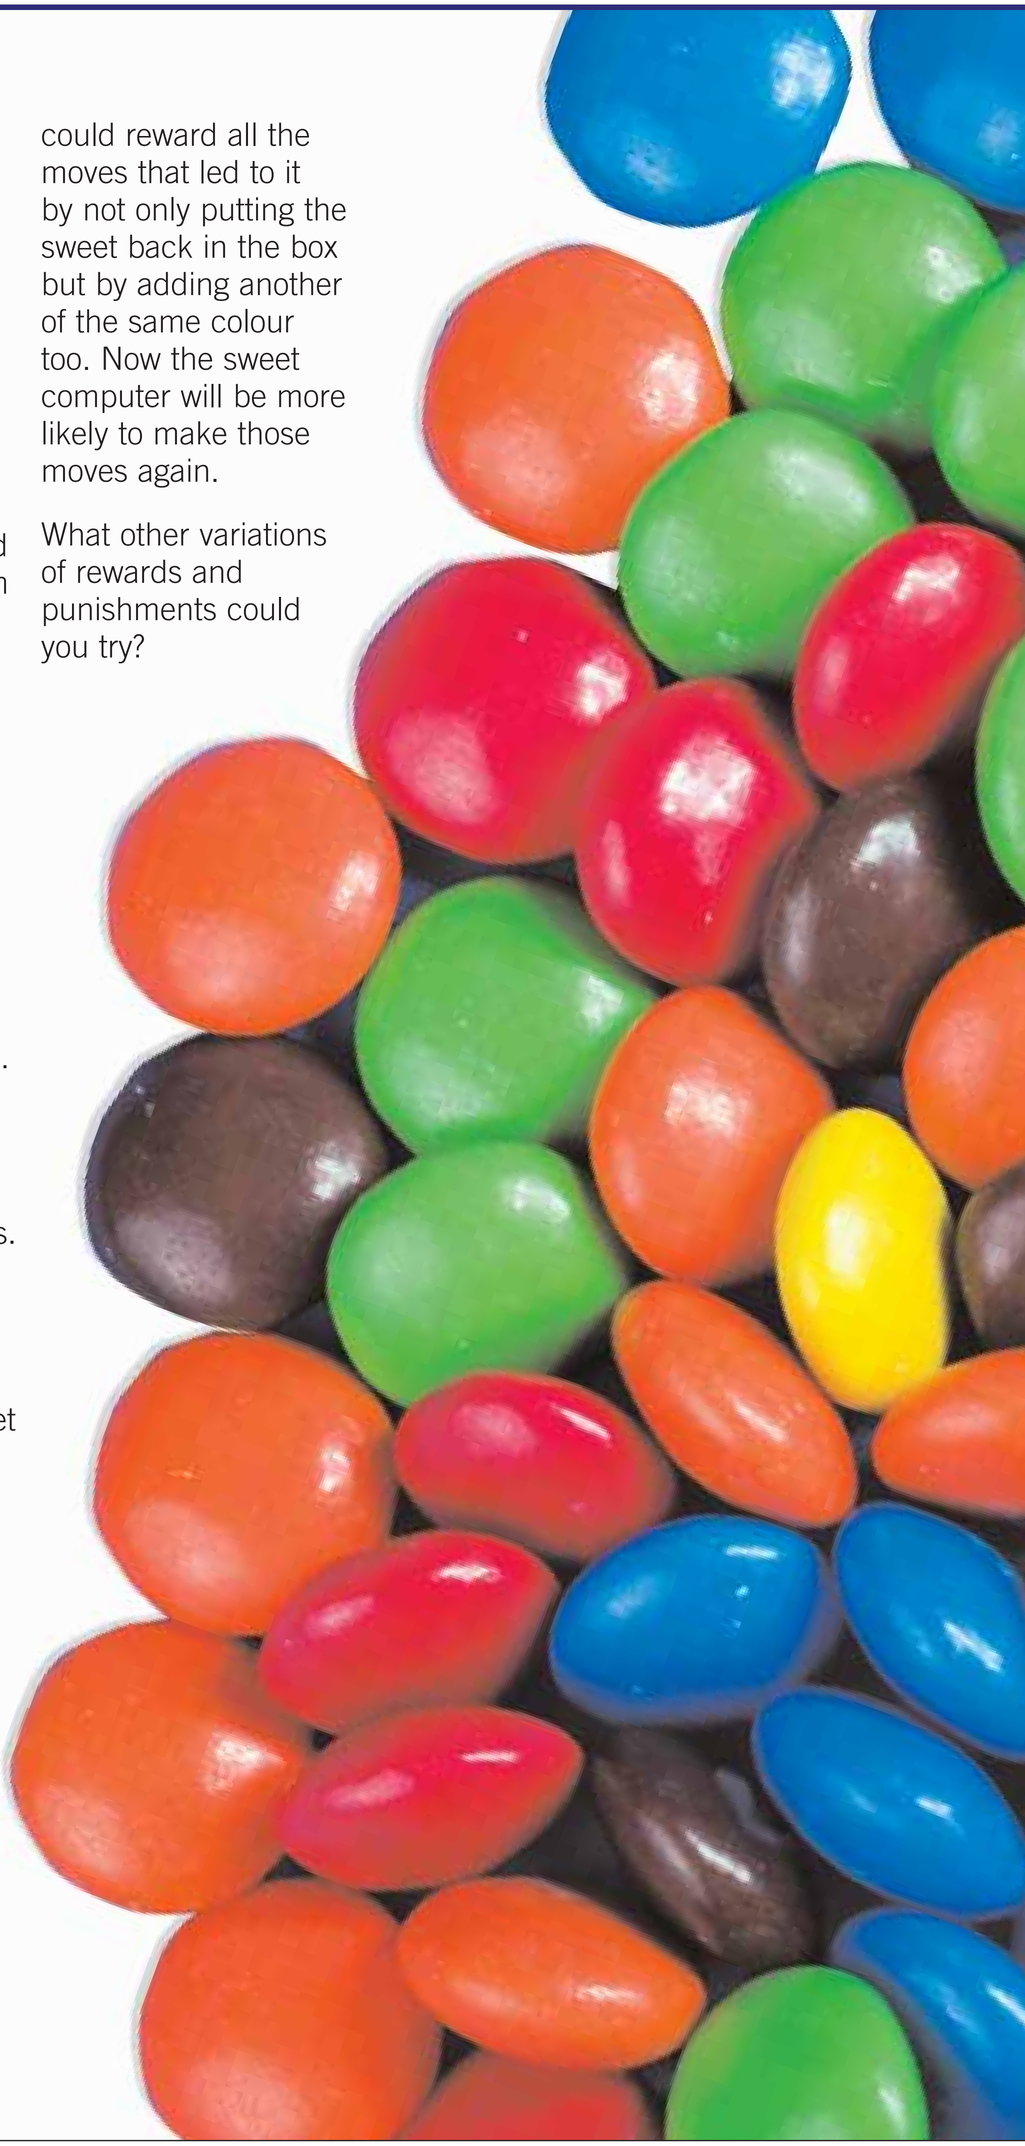

# Where's the understanding?

## One of the problems with accepting that a computer can be 'intelligent' is that it's just following rules written by a programmer. Isn't all the intelligence in the person who wrote the rules?

**Rules! Rules! Rules!**
The worry boils down to the idea that if you're just following rules blindly you don't have to understand anything, and if you don't understand then that isn't intelligence!

**Inside the Chinese room**
John Searle, a philosopher, came up with a thought experiment that captures this worry. Imagine you are locked in a room. Cards with Chinese writing on appear through a slot. You can't read Chinese. Instead you find the symbols on the card in a BIG book of replies. For every sequence of symbols you are given, it has another series of symbols – the correct Chinese answer to what came in. You copy it onto a card and push it out another slot.

Now you have no idea what any of it meant. However to the Chinese person asking you questions by putting cards in and reading your answers, you seem to understand Chinese perfectly. The Chinese person doesn't know you are just looking things up in a book.

You understand nothing but seem to understand everything!

**It's all in your head**
No one knows where understanding happens in our brains either. Is your head like a Chinese room with your brain just following complicated rules? If so eventually we will be able to make computers that do the same thing. They will then be able to understand in the same way we do.

**Losing consciousness**
What it boils down to is that we don't know what consciousness really is. Finding out is one of the great challenges for scientists to solve.

Consciousness isn't all or nothing though. What about those different animals we thought of back at the beginning of this book? What about a newborn baby, or a baby about to be born in a few days? What about you when you are asleep?

# creative computers

## There is more to being a human genius than being good at IQ tests. Being creative is just as important as being straight clever. A robot could never do great art, could it?

**Aaron can**
Aaron is a successful American painter. Aaron's paintings sell for thousands of dollars in the US. They have been displayed in art galleries worldwide including in London's Tate Modern Gallery and the San Francisco Museum of Modern Art. Oh and by the way, Aaron is a robot!

Yes, Aaron is a robot, controlled by artificial intelligence, and built by Professor Harold Cohen to be a creative machine. Aaron never paints the same picture twice. Once it starts it works by itself, with no human help. Aaron just draws and paints following the rules for art that it has been taught.

**Ed versus Zach**
In London, 2005, two computers took part in the ultimate cyber head-to-head. One computer ran an artist program called "Ed" by Ed Burton and the other a program called "Zach" written by Zach Leiberman.

They took it in turns to draw what they saw on each other's screen. They drew each other day after day generating wonderful interactive art.

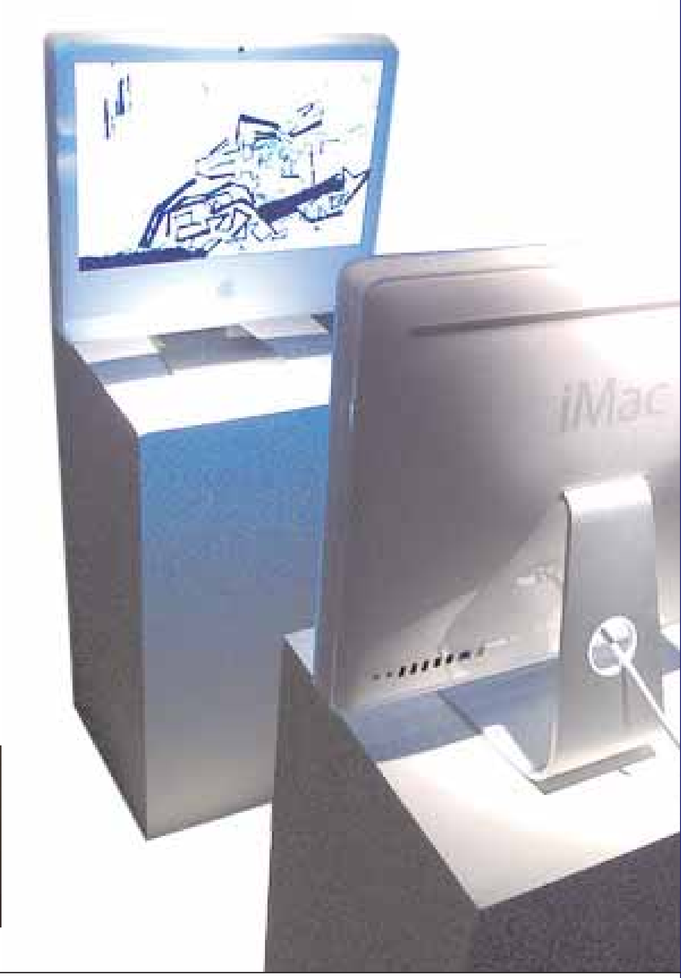

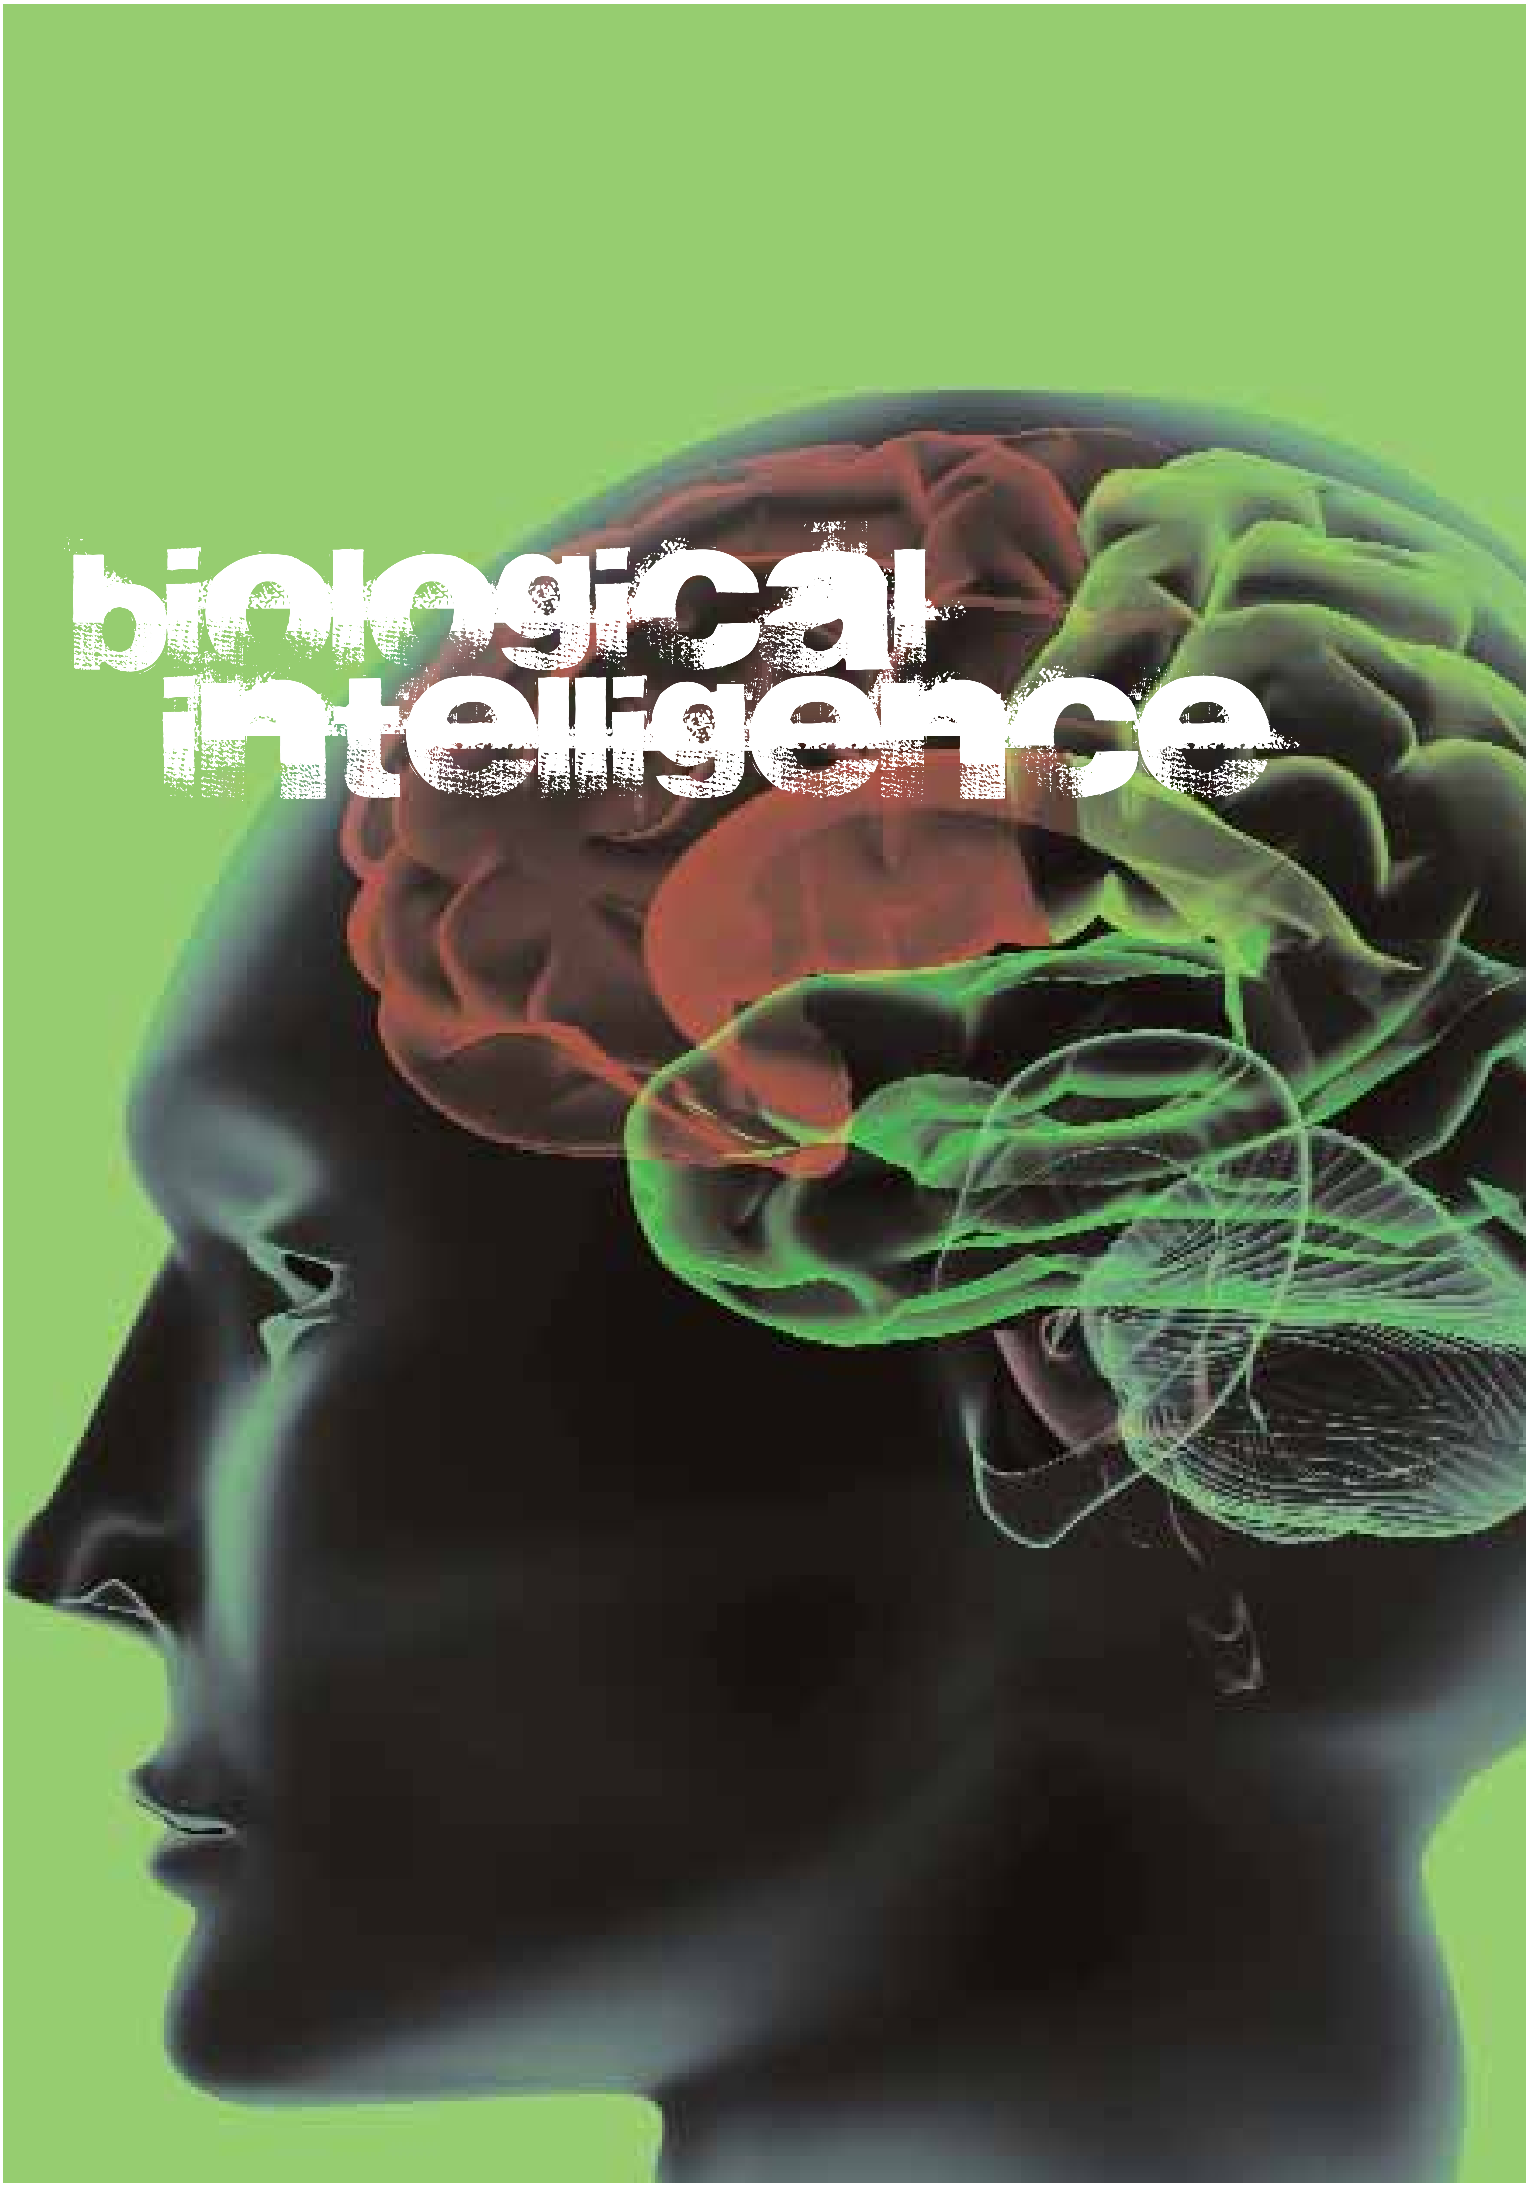
biological
intelligence

# Copy the machinery

We've seen that one way to give machines intelligence is to work out what we can do, and then try and write computer programs that can do the same things.

Another way is to work out how our brains actually work and copy the machinery. A computer that works in a similar way should then be able to do similar things to a human brain.

# it's just an illusion

## So how do our brains work? It turns out that your brain is following rules too! Not ones like "Keep off the grass", but rules even we don't know we're following.

One way scientists have shown our brains are following rules is using optical illusions. They happen in situations where our brains' rules don't fit – a bit like the piece of paper trying to play second at noughts and crosses.

**Stare at this eye. Perhaps wiggle the page a bit.**

Does it seem like the middle is floating and jiggling about? It's not of course! It's just a picture on a flat piece of paper. The pattern breaks rules that your brain uses to tell if something is separate from its background, so you are fooled by it!

Tricking your brain's rules can be really useful. Your TV is fooling you into thinking that the people, cars and animals on your screen are actually moving. But what's really on the screen is just a series of still pictures being shown one after the other. They're shown so quickly that your brain can't keep up, so it merges them together and gives the illusion of movement.

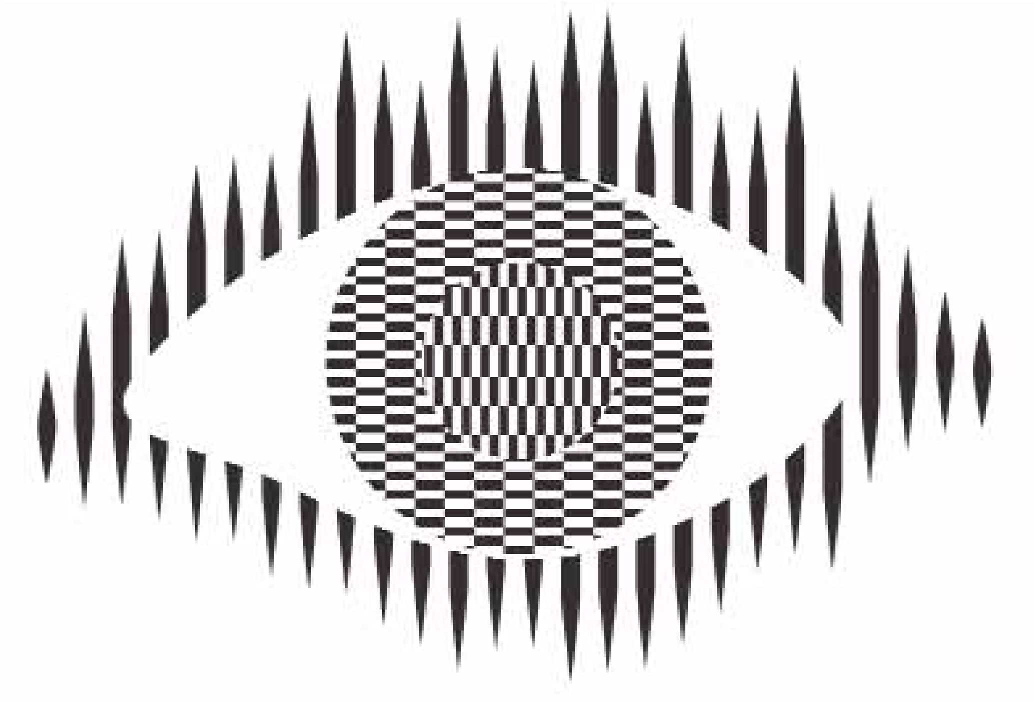

# neurons

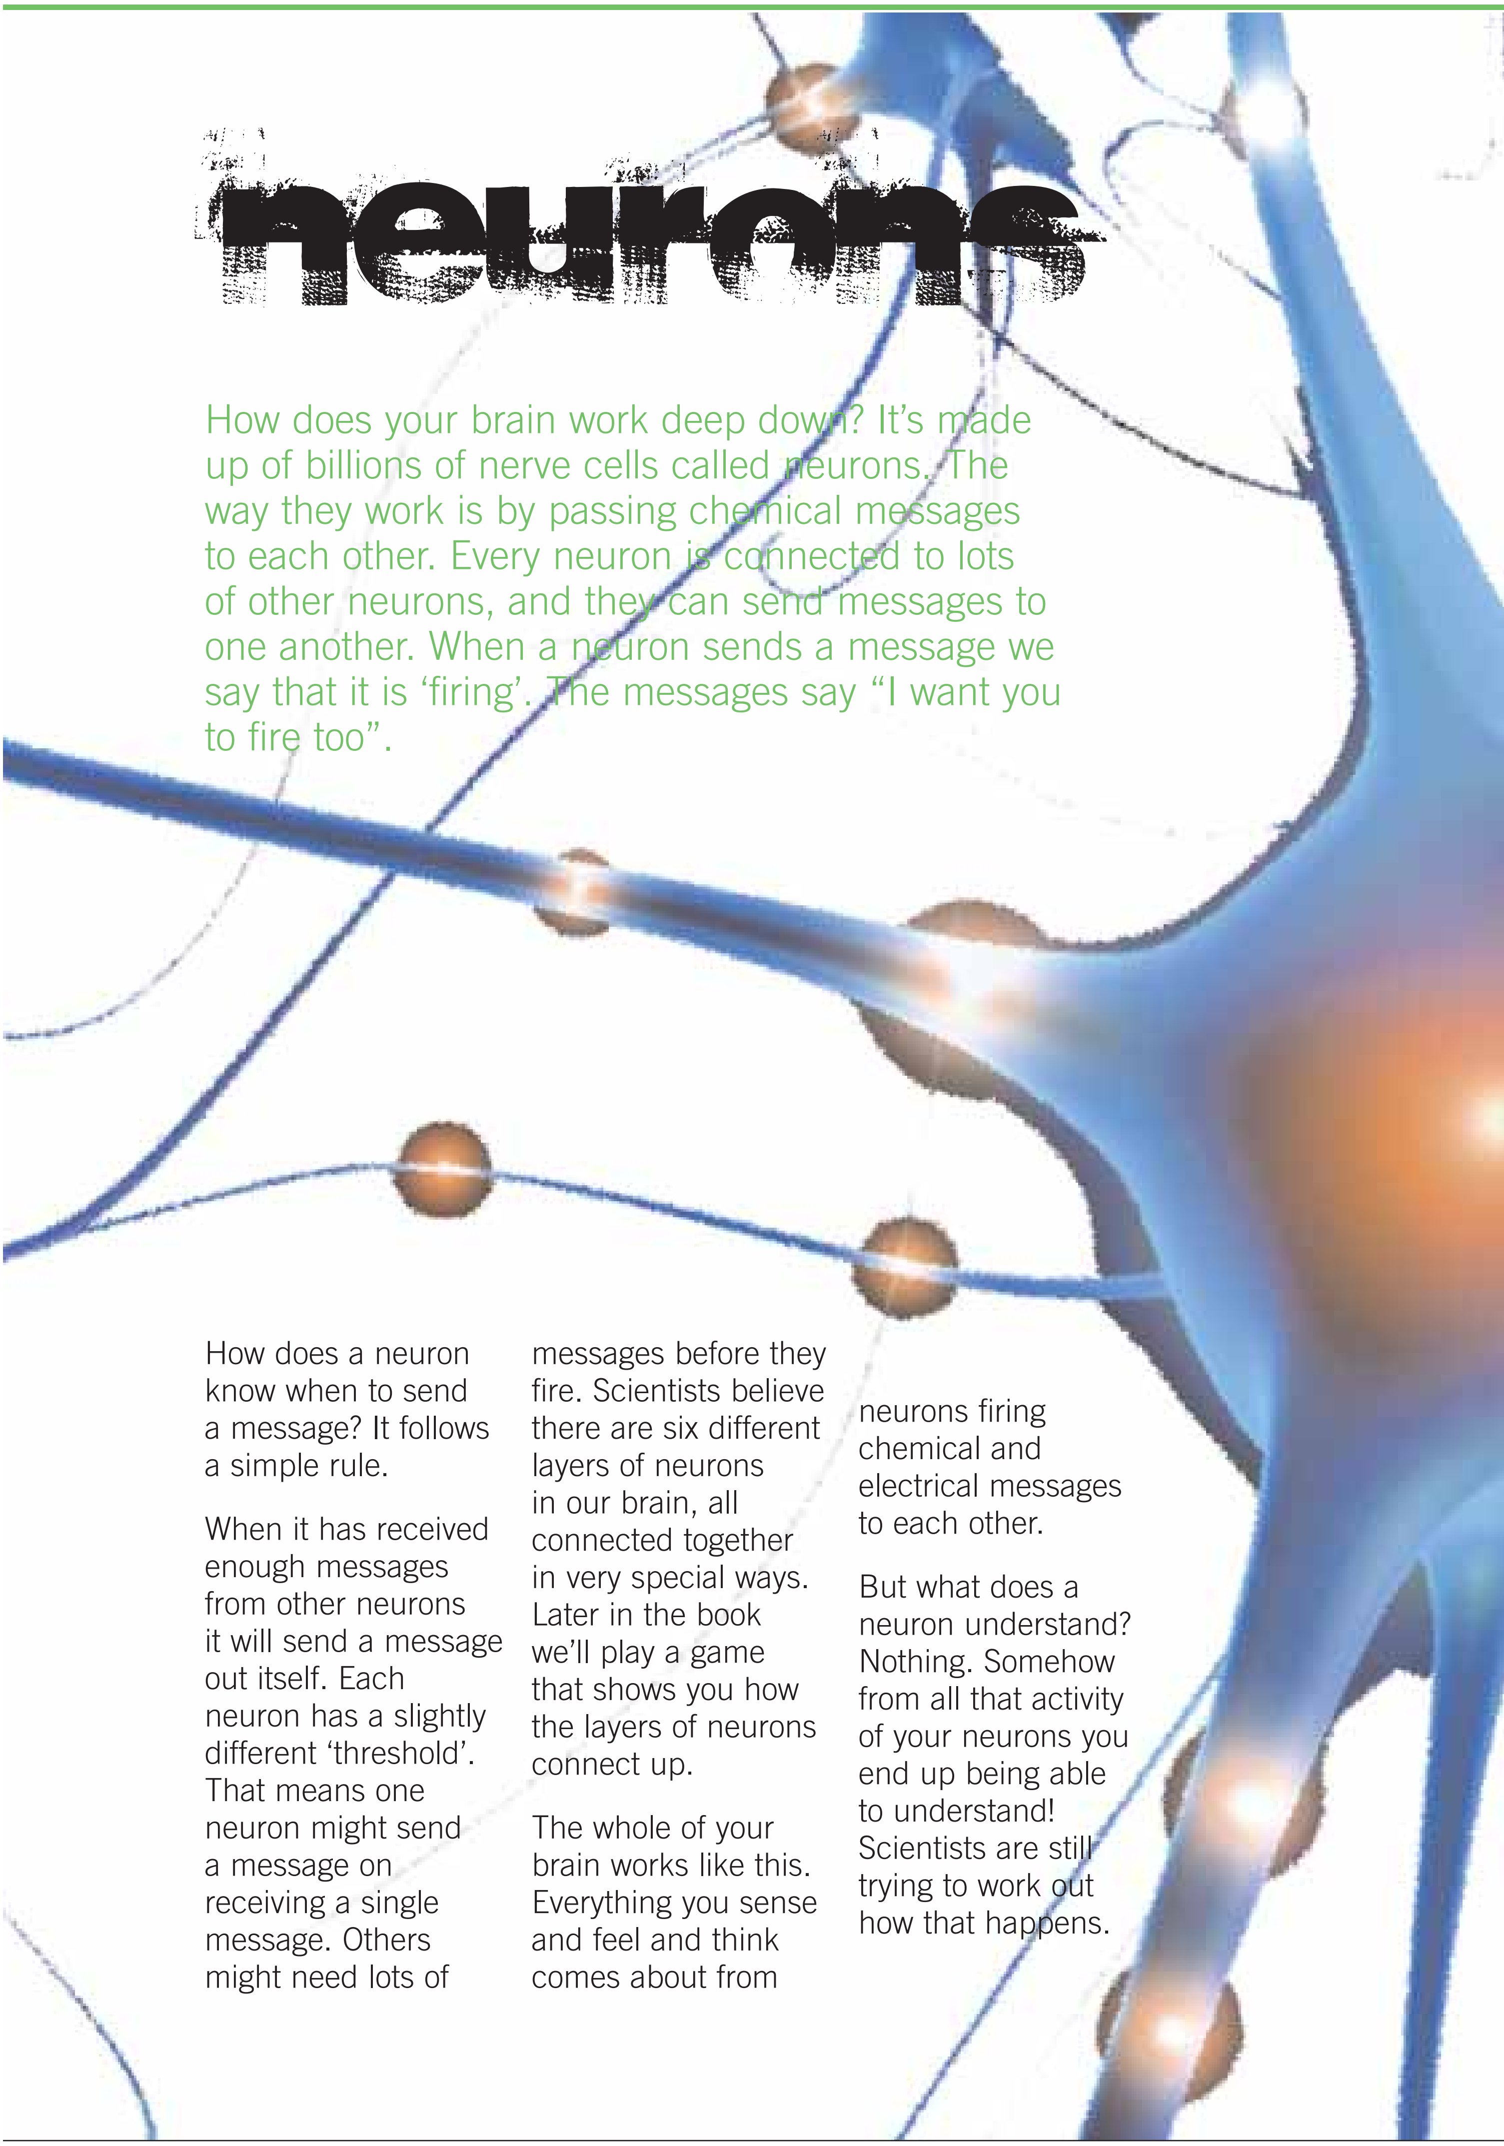

How does your brain work deep down? It's made up of billions of nerve cells called neurons. The way they work is by passing chemical messages to each other. Every neuron is connected to lots of other neurons, and they can send messages to one another. When a neuron sends a message we say that it is 'firing'. The messages say "I want you to fire too".

How does a neuron know when to send a message? It follows a simple rule.

When it has received enough messages from other neurons it will send a message out itself. Each neuron has a slightly different 'threshold'. That means one neuron might send a message on receiving a single message. Others might need lots of messages before they fire. Scientists believe there are six different layers of neurons in our brain, all connected together in very special ways. Later in the book we'll play a game that shows you how the layers of neurons connect up.

The whole of your brain works like this. Everything you sense and feel and think comes about from neurons firing chemical and electrical messages to each other.

But what does a neuron understand? Nothing. Somehow from all that activity of your neurons you end up being able to understand! Scientists are still trying to work out how that happens.

# build a brain to play snap

You can build a working brain of your own out of rope and toilet rolls. You will need seven friends to make it work. Alternatively you can make a small version out of strings and beads and operate it yourself, though that won't be so much fun.

First you need to build the hardware – actually in a brain it's more like squishyware.

Get some rope or string and cut two lengths about a metre long and two lengths that are 2 metres long. Tie each short piece to the middle of a long piece and lay them on the floor in two Y shapes.

Each friend will control a neuron in your brain. They provide the energy to allow it to work. Get them to stand in the positions shown in the diagram: one person holding each knot and the others holding the ends. A single person should be holding the bottom of both 'Y's. Pull the whole contraption tight.

OK, so you have artificial neurons connected together, but neurons can only do things if they have chemicals to deliver messages.

There are 100 billion neurons in a human brain.

Our chemicals are going to be toilet rolls. Thread a toilet roll on the end of the arms of each Y. Next thread one on the end of each base. Move those last ones up to the knot.

**Turn the page to see how to make your brain work.**

**How many neurons in their brains?**

| | |
|---|---|
| octopus | 300 million |
| rat | 21 million |
| snail | 11,000 |
| ant | 10,000 |
| slug | 7,000 |

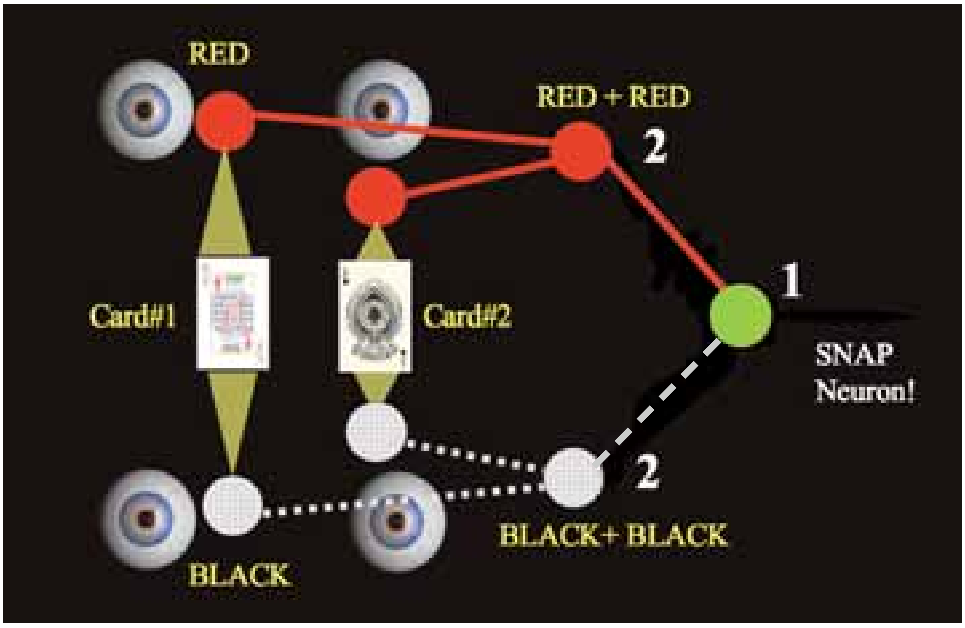


A fruit fly gets by with only about 100,000 neurons. So why is it so hard to swat? It makes clever use of the few neurons it does have.

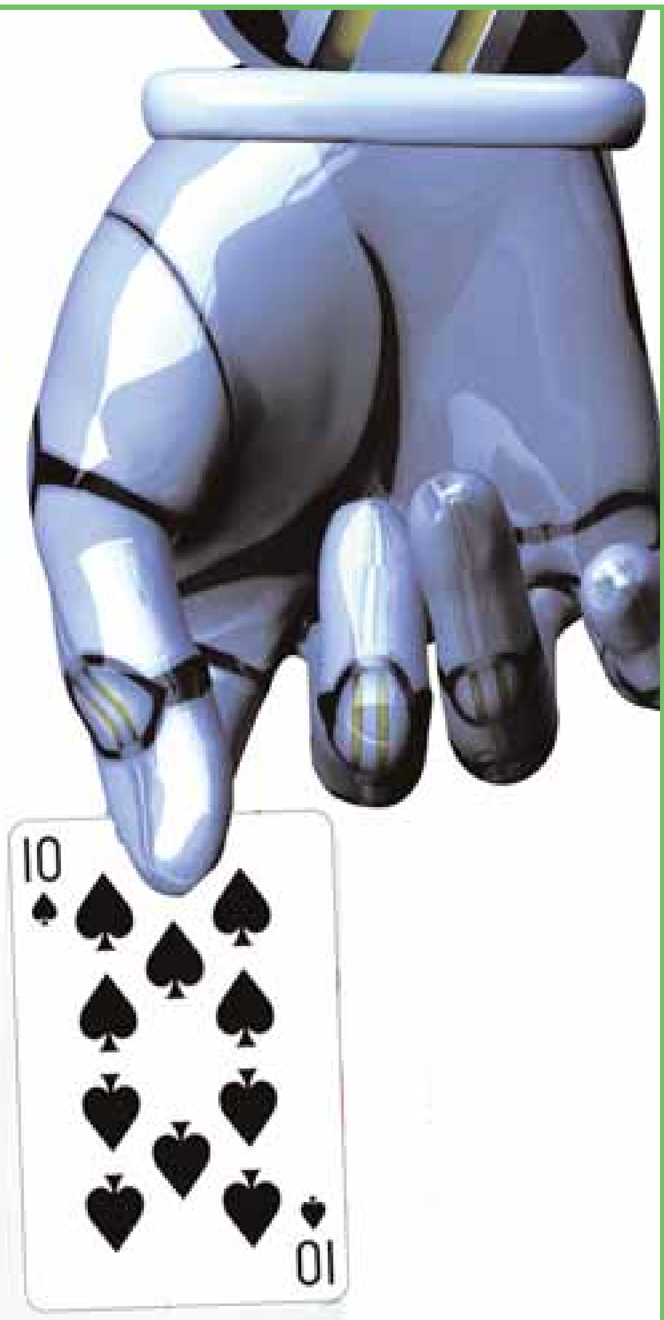

### Tell me when to fire!
Now you need to bring your brain to life. The neurons need rules!

The rules in a neuron tell it when to fire. For our artificial brain that means a rule to tell your friends about when they should do an action – usually sending a toilet roll down the rope to the next person.

We're going to play a simpler version of snap – only the colours have to match, not the numbers.

### Senses
The four people at the ends of the Ys are EYE neurons. Brains get information from the outside world through the body's senses. Our brain has eyes to sense the world and the eye sends messages to the rest of the brain through neurons. Our artificial EYE neurons each have a different rule. Each is focusing on one thing that can be seen. Cards will be placed on the table in one of two positions: position 1 or position 2.

- The first EYE neuron fires if there is a Red card in position 1.
- The second EYE neuron fires if there is a Red card in position 2.
- The third EYE neuron fires if there is a Black card in position 1.
- The fourth EYE neuron fires if there is a Black card in position 2.

To fire, all the neurons have to do is shoot the toilet roll along the rope.

### Deep in the brain
The two people holding the knots are INTERNEURONs. Interneurons are further down in the brain. Unlike the eye neurons they aren't connected to the brain's senses, only to other neurons. Their rule in our brain is for them to fire their toilet roll only if they get two toilet rolls fired at them.

### Shout!
The final person who should be holding two ropes is a SNAP neuron. A SNAP neuron is different. When it fires it controls the voice connected to our brain. When the SNAP neuron fires, the person just shouts "SNAP" as loud as they can. When does it fire? It fires if a single toilet roll arrives.

### Play the game

Your artificial brain is now ready to play snap! The EYE neurons should watch their own spot on the table. Someone else should lay down cards in pairs. If both are red or both are black our artificial brain should end up shouting “SNAP!” If they are different colours it should stay quiet – just like when you play snap without an artificial brain.

After each round of two cards, the brain needs to reset. In a real brain all the chemical messengers have to wash back to where they came from. In our artificial brain it’s the toilet rolls that have to go back.

Make up two artificial brains and you could have a real game of snap between them. Alternatively you can make small brains with straws and beads. Then it’s your job to push the right beads about, and you can only shout snap if a bead gets to the SNAP neuron!

May the fastest brain win!

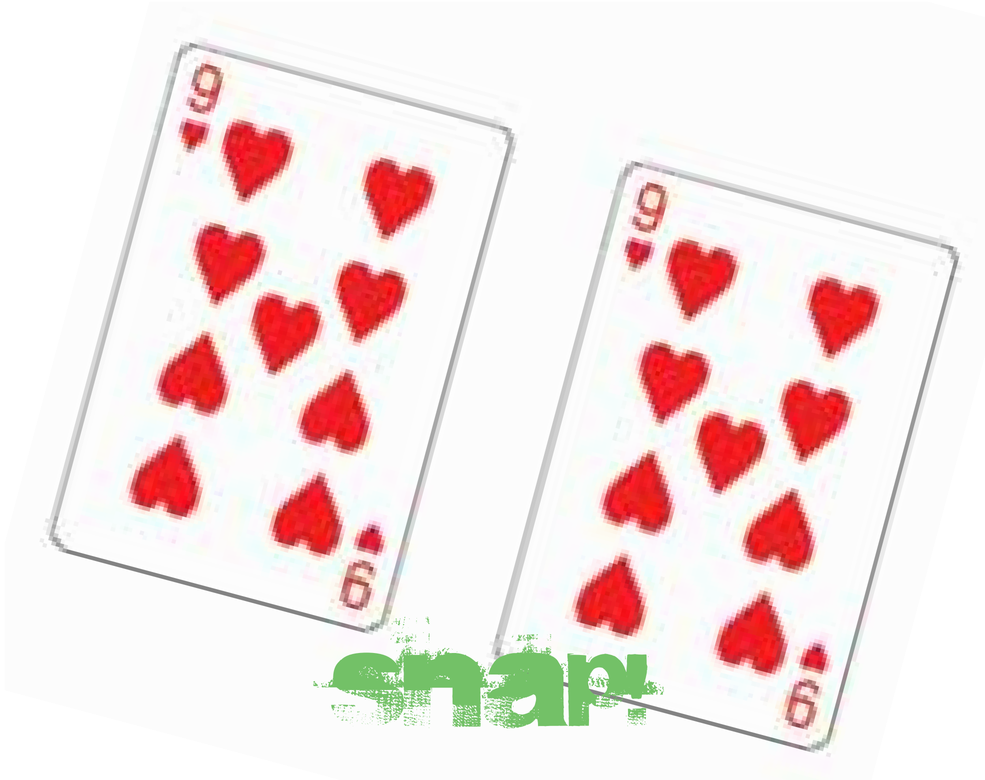


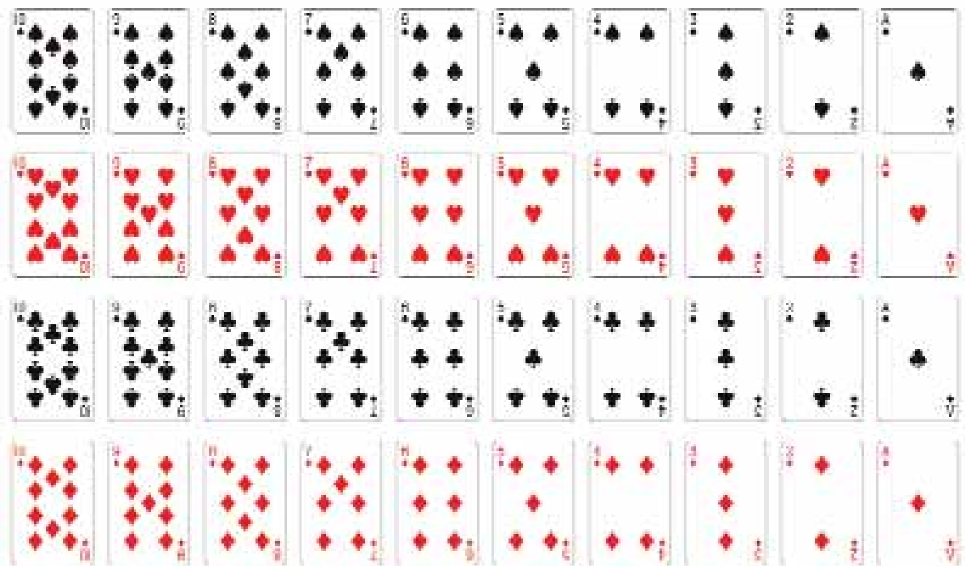

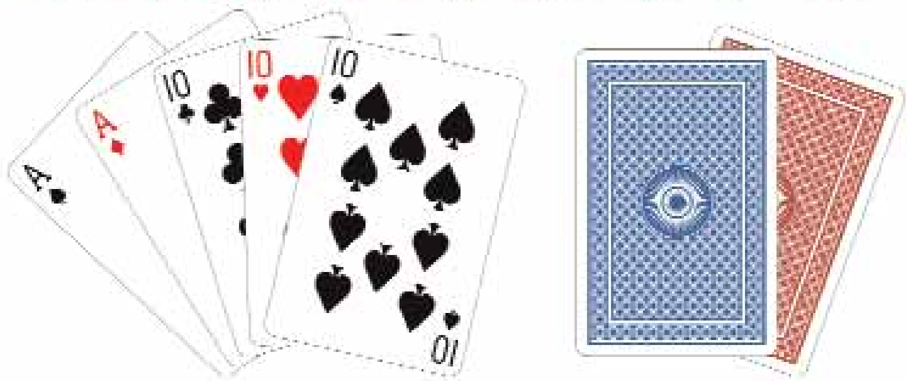

## Spot the tiger!

Why not try and make a brain to do other things, like identifying animals? Let's try to make a brain that recognises whether it's looking at a tiger, which has stripes, paws and sharp teeth, rather than a zebra, which has stripes too but hooves instead of paws and no sharp teeth.

This can get a bit complex, so you might want to experiment with beads for toilet rolls and straws instead of rope before building a full-size brain.

You need two MOUTH neurons – one to shout "ZEBRA", the other to shout "TIGER". You will need one EYE neuron that fires if it sees STRIPES. A second EYE neuron will fire if it sees "SHARP TEETH". A third eye neuron will fire if it sees "HOOVES".

Now you just need to connect them in the right way.

The EYE neurons may have two lines coming out of them, each with a bead to fire. One might go into the tiger-spotting part of the brain and the other into the zebra-spotting part of the brain.

When you need a bit of brain that fires when two things are true – it has stripes AND it has sharp teeth – you just need a Y-shape like in our snap brain. The middle INTERNEURON of the Y has a rule to fire if beads arrive on both the lines coming in.

When you need a bit of brain that can tell if one thing is true or another thing is true, you just have a neuron that fires as soon as it gets any bead on its lines coming in. That's how the snap MOUTH neuron worked as it had to fire if there were 2 red or 2 black cards.

Think up some other similar animals and see if you can create brains that tell them apart. How about one that can identify whether it is a bird, fish, mammal, insect or reptile? But don't forget about the awkward ones like bats – a mammal that flies, or the duck-billed platypus – a mammal that lays eggs.

# neural networks

It turns out you can write a computer program to act like a neuron. It just has a way of following simple rules of firing. By putting together lots of neuron programs into a single big program and giving them a way to send messages to each other (like the toilet rolls), you can create a program that works just like our artificial brain did but inside a computer. The computer then has a brain that works much like our real brains do.

The picture on the next page shows a robot with a brain just like that. Called Blade, it was made out of Lego by Zabair Qureshi, a student at Queen Mary, University of London. Rather than detecting animals, Blade was able to hear the tone of someone's voice. It can tell the difference between a loud, angry voice and a soft, calm voice. The rules in its neural network, rather than making the robot shout a word like in snap, changed its expression. Blade could look happy, sad or surprised, depending on what it heard.

A difference to the artificial brain we made out of rope is that the neurons in Blade's brain learned their rules by experience. Zabair and his friends shouted and talked softly to Blade a lot, training it how to react – a little like you might train a dog. Over time it learned to look sad when Zabair was angry and to look happy when he sounded calm.

When Zabair finally had to take Blade apart, he said it was difficult to do. He'd got so used to Blade's human-like reactions it felt like it was much more than just a pile of Lego bricks.

Your brain contains about the same number of nerve cells as there are raindrops in an Olympic size swimming pool.

'Moore's Law' says that every year and a half the cost of electronic chips drops by half, and their power doubles. Some researchers believe that by 2019 even a home computer will be able to do the same number of calculations as the human brain.

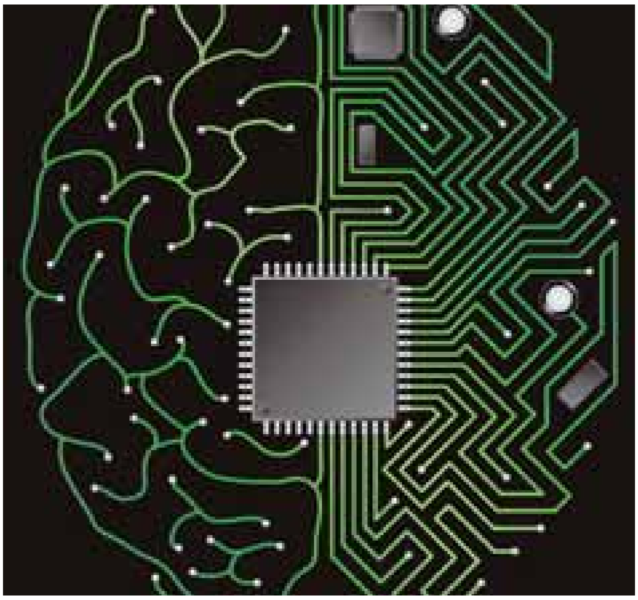

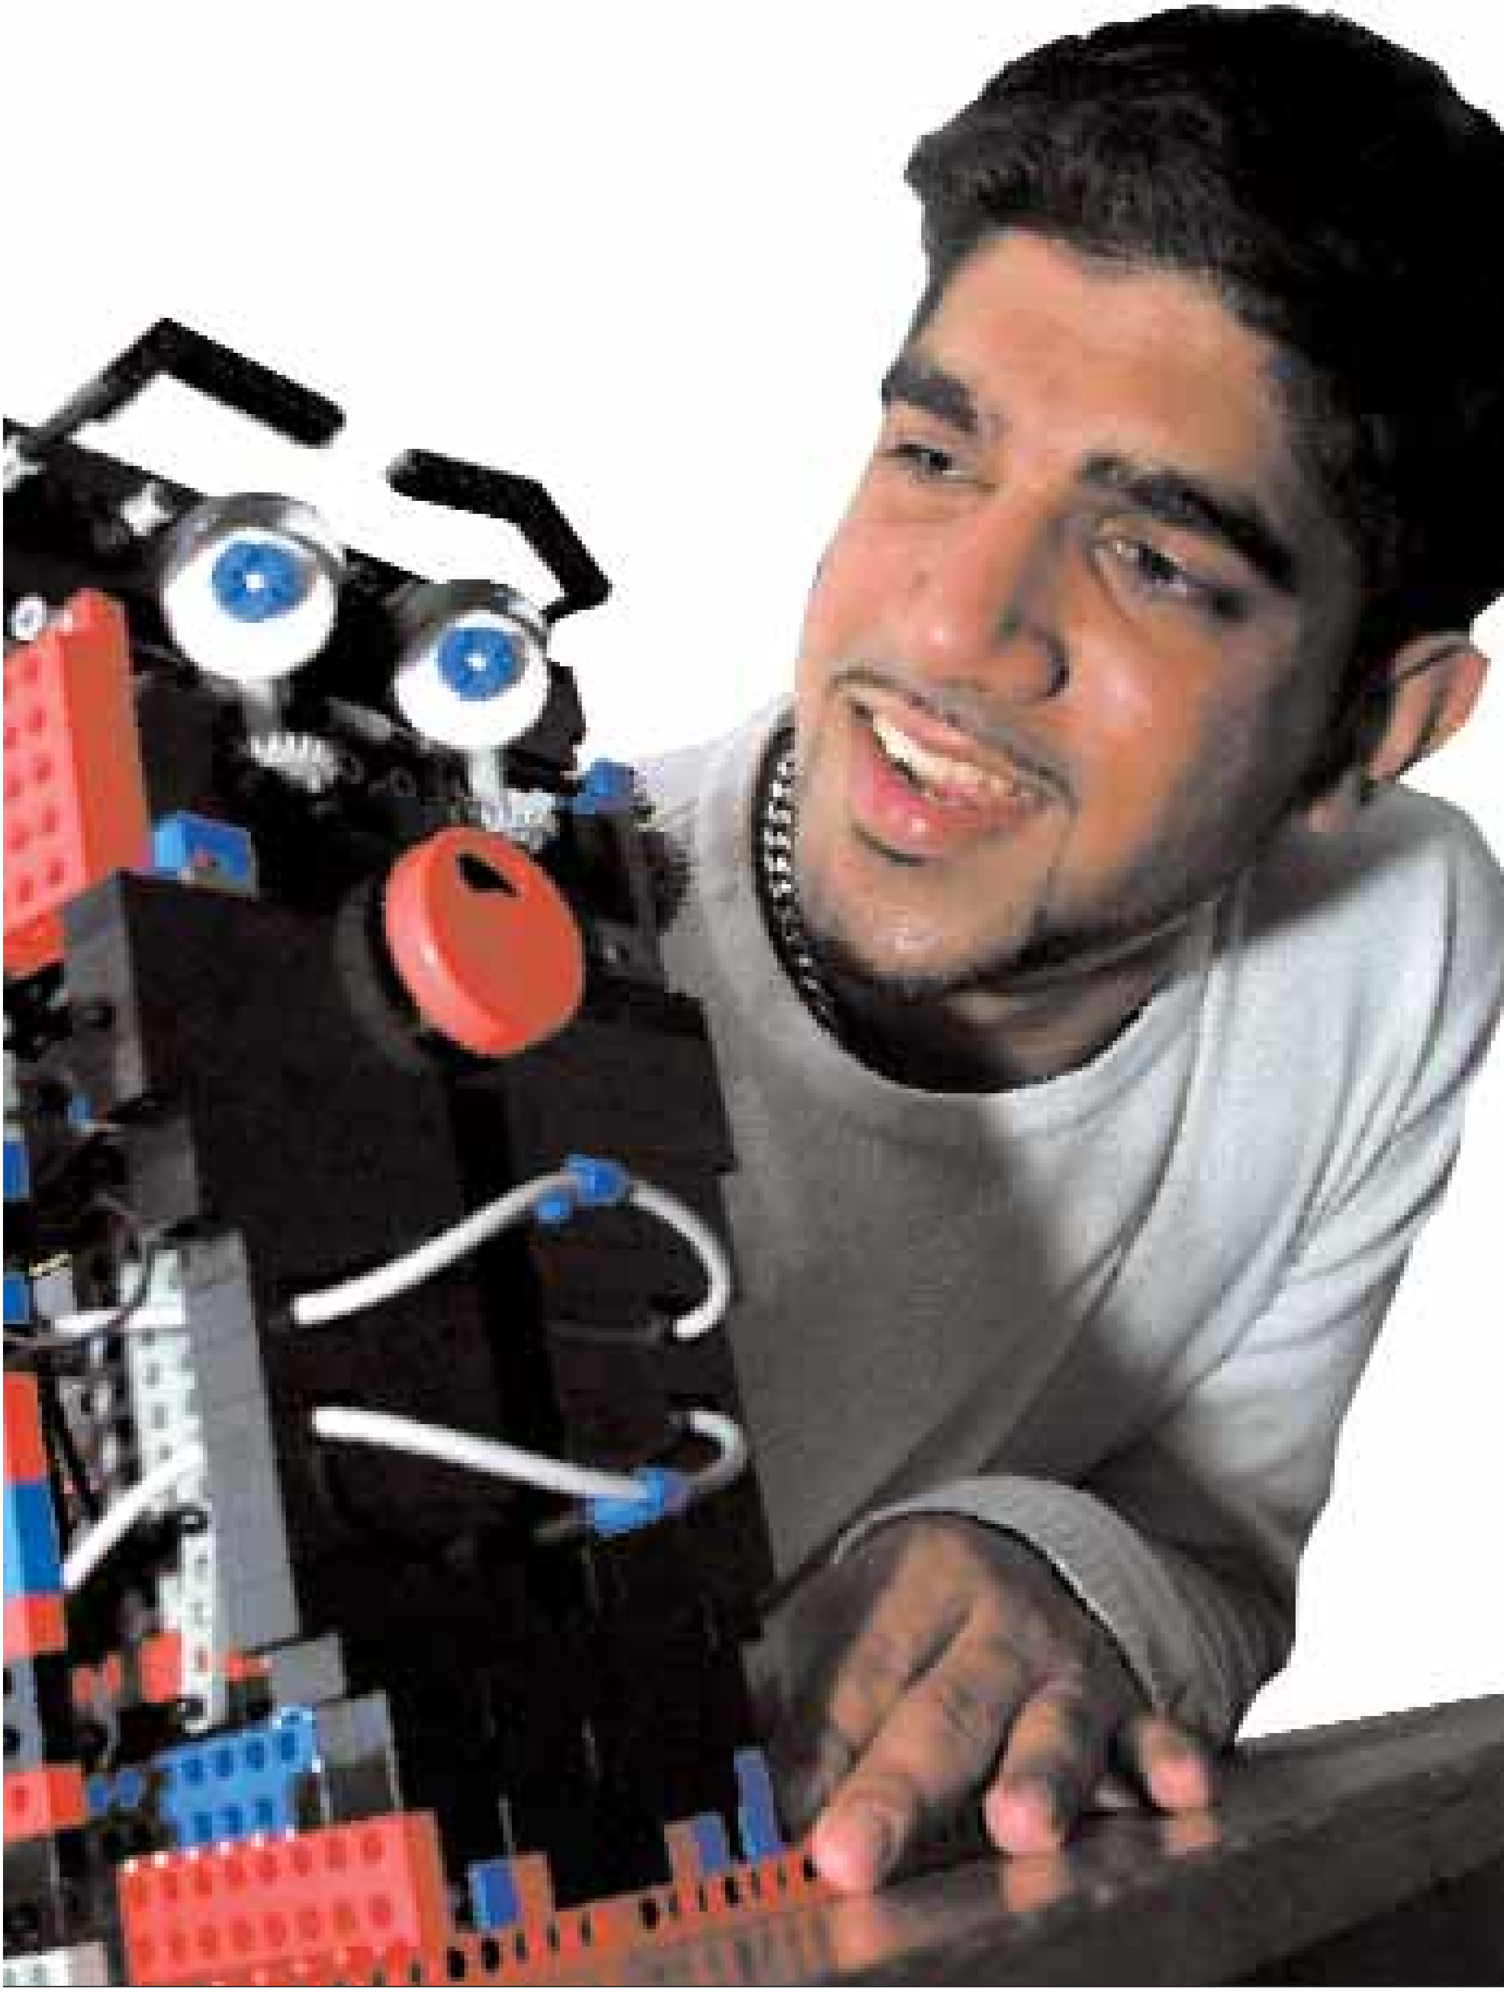

# robots

# intelligent robots

Put an intelligent computer into a mechanical body and you get a robot. Science fiction films are full of them, but will robots like R2-D2 and C-3PO ever really exist?

**Czech robots**
'Robot' comes from the Czech word 'robota'. It means 'compulsory worker'.

The word 'robot' was first used in 1921 by Czech writer Karel Čapek in a play called RUR (Rossuum's Universal Robots).

**Gastrobot**
In 1996 two America scientists created Gastrobot, a robot that digests real food. Its stomach turns the food into carbon dioxide, which is then used to power the robot. Burp!

**Climbing stairs**
The car company Honda don't just make cars. They're also trying to build humanoid robots like C3PO in Star Wars too. In 1996 they revealed the P3. It could walk by itself, and even climb up stairs, which is more than a Dalek can do!

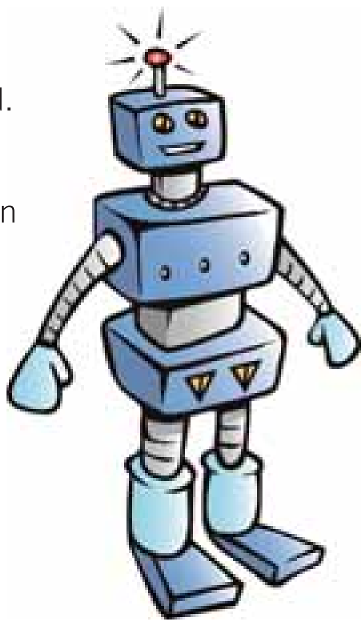

**Ring that bell!**
In 2002 Honda's new ASIMO robot rang the opening bell at the New York Stock Exchange. ASIMO can do even more than P3; it can dance, run, play musical instruments and even conduct an orchestra.

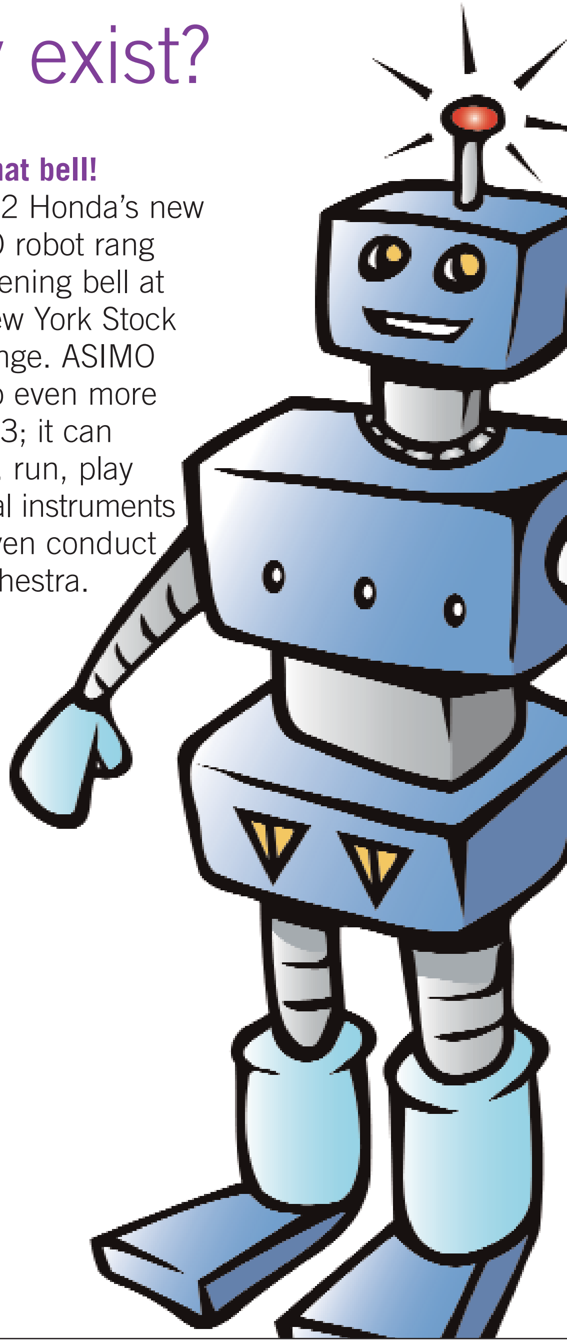

# making it happen

We've made brains out of rope and toilet rolls, but how can a brain actually turn the things it senses into useful action? The earliest creatures with eyes could probably only sense the difference between darkness and bright light. How useful would a really simple 'eye' like that be to a creature?

Let's do a thought experiment. Let's imagine creating a 'wheelybot'. It is a robot that needs strong light to survive. Perhaps it's solar-powered. That means it needs to stay in the light – it wants to move towards light and away from the dark.

How complicated a brain would our wheelybot need to have to always move towards the light?

Suppose the wheelybot has a wheel connected to a motor with two terminals to connect wires to. If a signal comes in to the first terminal (the FORWARDS terminal), the wheel turns forwards so the wheelybot moves forwards. If a signal arrives at the second terminal (the BACKWARDS terminal) it makes the wheel turn backwards, and the wheelybot moves backwards.

Now suppose our eye is connected to a neural network like the one we made out of rope and toilet rolls. This neural network has two EYE neurons coming out of it. The first fires if there is bright light ahead. The second fires if there is only darkness ahead. Otherwise nothing fires. All we have to do is connect its bright EYE neuron straight to the FORWARDS terminal and the other to the BACKWARDS terminal. Our imaginary wheelybot will be drawn towards bright light, just like a moth, but will scuttle away if it's dark.

Now think again. What would happen if we connected the eye neurons to the motor terminals the other way round? With a simple change to the wiring the wheelybot's behaviour has changed. Now it slinks away from the light and into dark places – just what it might need if there was a predator out to gobble it up if it came out in the daylight.

The way neurons in brains are connected together makes a real difference to how creatures behave. As we've just seen, even a simple swap can lead to big changes.

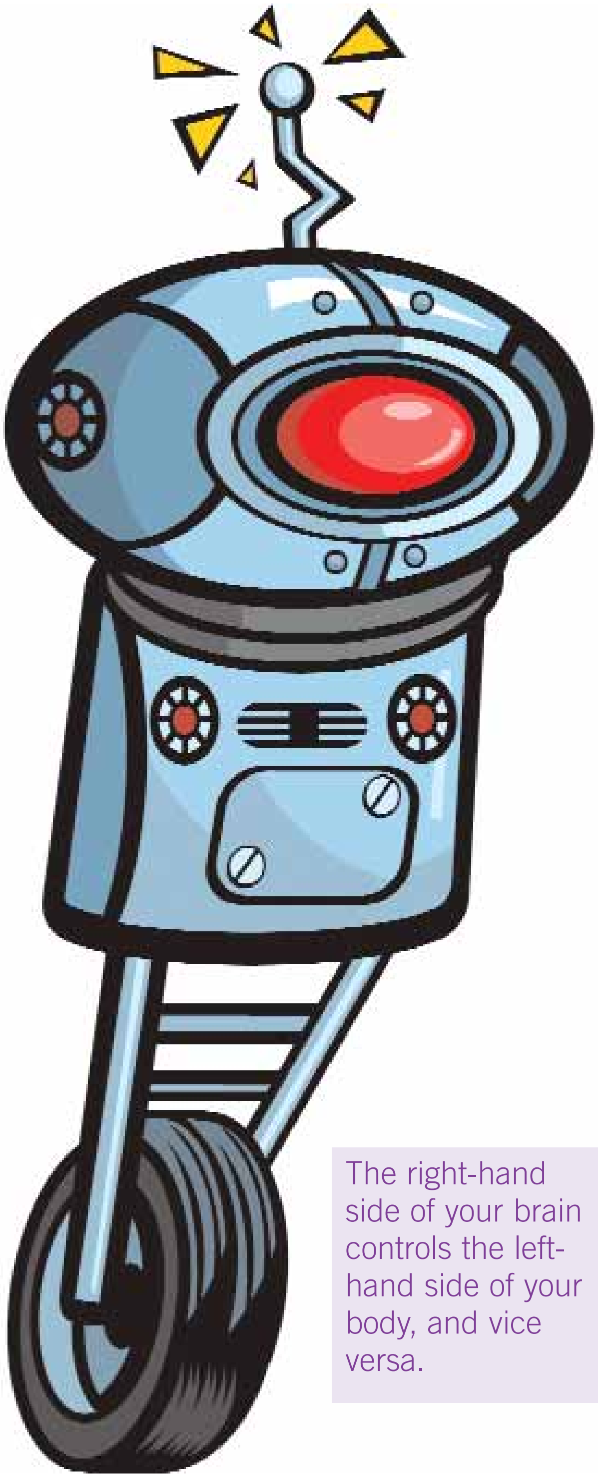

The right-hand side of your brain controls the left-hand side of your body, and vice versa.

# polite robots

Robots already work for us. They build cars in factories and help repair satellites in space. If robots are going to be useful at home too they will need to learn good manners. After all, intelligence is about being able to get on with others too.

**Did you say thanks?**
Saying please and thank you, not standing too close to someone you don't know very well, and asking "Excuse me, can I get past? are all examples of something called 'social etiquette' (a posh word for good manners). It's a set of rules that lets people in a group get on with each other. To fit in, any robots in the group will have to understand the rules too!

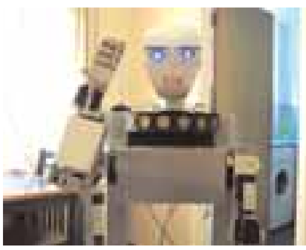

That's where computer scientist Kerstin Dautenhahn comes in. Her robots have a special kind of intelligence: social intelligence. They know how to behave properly around people.

Kerstin started out as a biologist looking at how stick insects' legs move.

Kerstin realised that studying biology was a great way to find ideas to help build robots, and that the social intelligence animals use to live with each other can help us to create polite, sociable robots.

She is trying to answer questions like "What's the best way for a robot to interrupt if you are reading a comic? Should it gesture with its arms, blink its lights or make a sound?"

What do you think it should do? (We'll tell you what Kerstin's group came up with on the next page.)

# friends at home

## Kerstin's group created a special flat to look at how people and robots got on together at home.

People lived in the house with the robots for weeks at a time.

### In your face

One experiment was about discovering the best way for a robot to approach a person. At first Kerstin's group thought that it would be best if the robot came straight towards the person from the front. People found that was too aggressive. Instead the robots are now trained to come up gently from the side.

### Comfort zone

The people in the house were also given special 'comfort buttons' to press if the robots made them feel uncomfortable. Quite a lot of people were happy for the robots to be close to them. In fact they were happy for robots to come closer to them than people they didn't know well.

This is because people see the robots as machines. As Kirsten says: "You are happy to move close to your fridge, and it's the same for robots".

### Friendly future

By building robots with manners and good social skills Kerstin and a team of people across Europe hope to create robots that will be our friends, helping to make our future friendly.

An 'android' is a robot that looks like a human. Researchers in Japan have found that people find existing androids creepy because they look like a human but don't behave like a proper human. People seem to prefer being around robots that look like robots.

# pet robots

## Make your own

In 1989 Dr Seymour Papert become the first Lego Professor of Learning Research at a top American university.

His wages were paid by Lego, the toy company. He helped them create new electronic bricks to be part of a kit for making toy robots.

The resulting toy was called Lego Mindstorm after a research paper that Dr Papert wrote. You can build your own Mindstorm robots using the bricks, then program them to behave the way you want them to.

Blade the emotional robot (see page 34) was made of Lego Mindstorm

## Furbys

Furbys came out in 1998, and were one of the first electronic pets you could buy in the shops. They react with sounds and facial expressions when you play with them. They know over 800 phrases to use to talk to their owner. They also have their own language called 'Furbish'.

## Sit!

Electronics company Sony have created a range of robotic dogs called the AIBO. These cute robotic pets can sit and walk like a real dog. They can even 'see' using a tiny camera built into their head. Using this camera AIBOs can play games with a brightly coloured ball. They are also able to learn and remember tricks their owner teaches them.

## Rabbit! Rabbit! Rabbit!

A Nabaztag is a clever electronic rabbit. It collects information from the Internet for you. It can read you weather forecasts, tell you the news, let you know the time, wake you up, download and play MP3s and read you email messages from friends…for starters. Nabaztags can even talk to fellow bunnies over the net and get married!

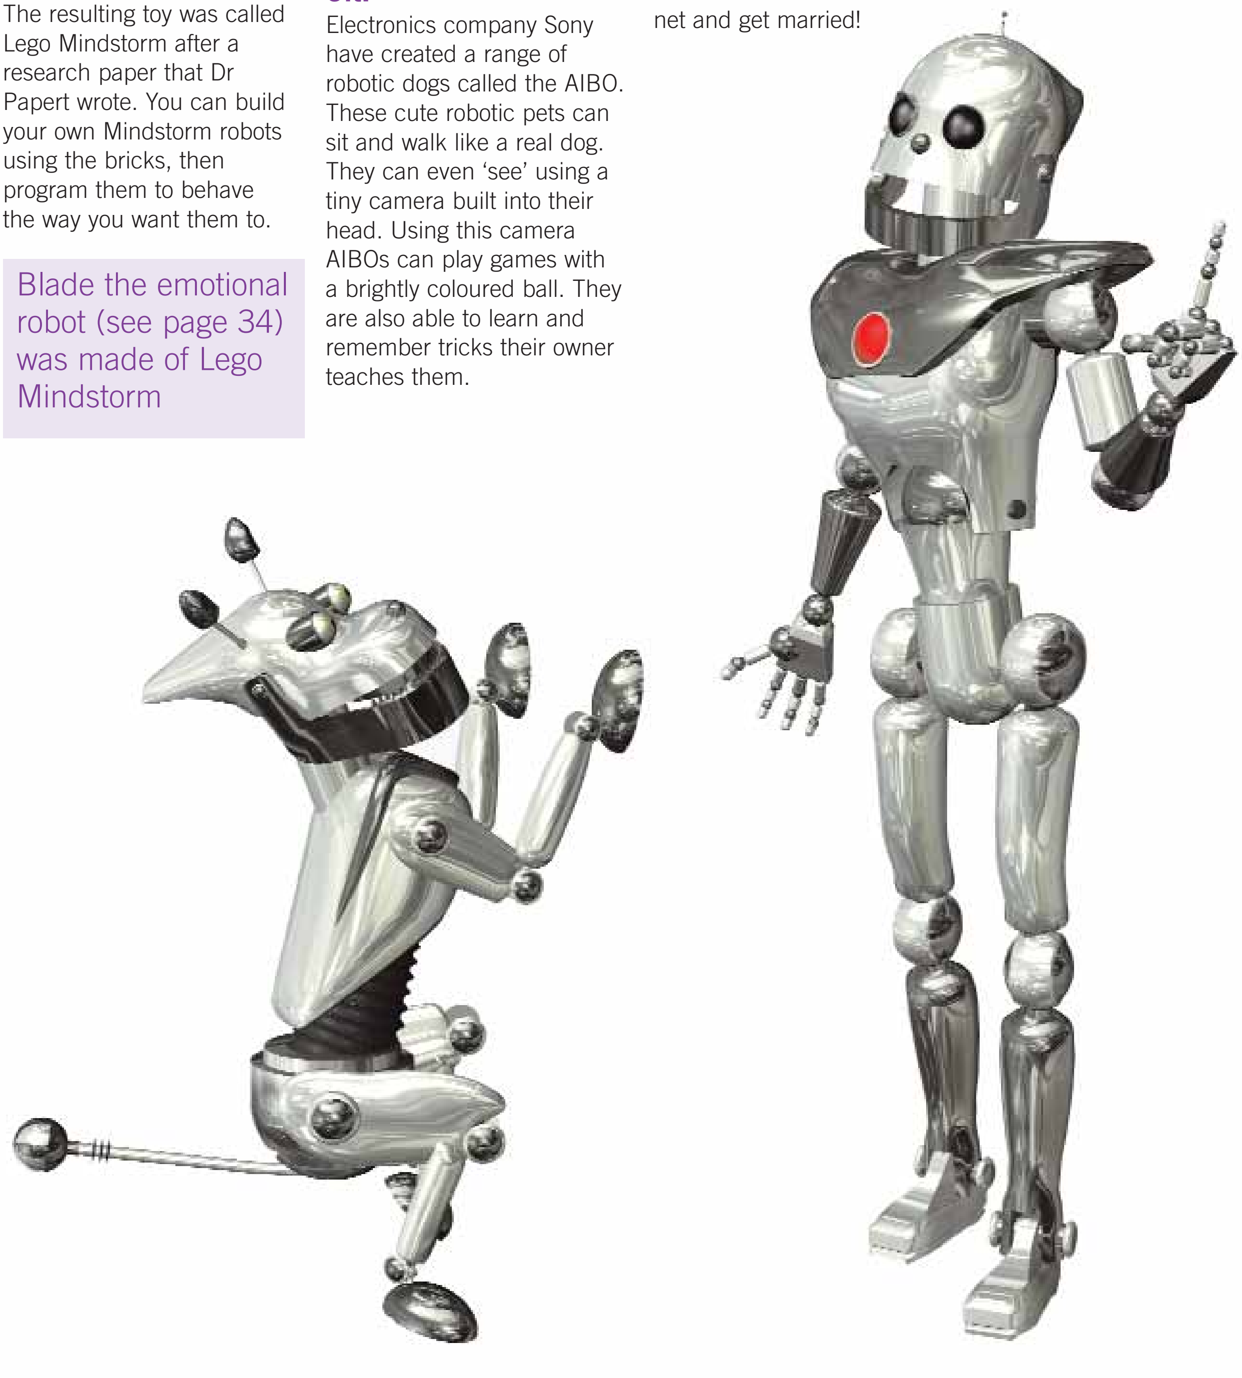

# the future

# human race?

## Who can create the fastest creatures to race over digital racetracks? Have a play with some artificial intelligence and find out at Sodarace (www.sodarace.net)

Sodarace is a free online Olympics where humans can try to win against artificial intelligence. Using your mouse to draw a springy creature, you can then add muscles to make it move. Muscles move back and forth in time to an energy wave that passes through them.  You can create cars or caterpillars, birds or bouncing boxes. It's up to you and your imagination. You can also see what other players have created and use their ideas to help create your own. You can then create racetracks, from flat and fast to steep and scary. All you need to do is point your mouse and click away. How will your creature do in the races? Will it make it to the finishing flag?

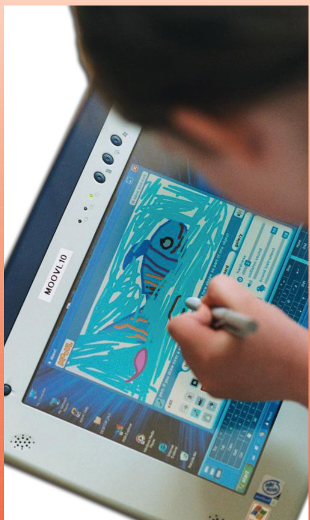

You can play with some artificial intelligence yourself too. On the website you'll find software that lets you create amoebas, which in the game are rolling wheel-like creatures. You can then ask the artificial intelligence for help to make them go faster.

### Learning by trial and error

The artificial intelligence works in much the same way as you might. It makes tweaks to the creature, and sees if that change makes it go faster. If it does then the artificial intelligence keeps the change. If the creature is slower then the artificial intelligence throws the change away. It's a bit like a human designer using trial and error to make things better. The artificial intelligence can try hundreds and hundreds of changes though. Each time it just keeps the best.

Is this really intelligent? Well, humans learn from their mistakes all the time. If something doesn't work you tend not to do it again, and that seems quite a clever thing to do. So is our artificial intelligence smart because it learns from its mistakes too?

Why not try and answer the question for yourself on Sodarace and join forces with the artificial intelligence to see if your team can be a winner.

Moovl, which is a cartoon version of Sodaconstructor, uses some of the same software as Sodarace. Computer scientists often do this: reusing parts means you don't need to waste time getting them to work from scratch.

### Sodarace: round one to humanity

The first sodarace to pit humans against artificial intelligence was won by a human, but only at the very last minute. An AI had designed what was the fastest creature on sodarace at the time, and its creators were all set to take the news of their program's intelligence to the public. But just before they did, a schoolkid in Canada did a bit of work on the AI-created creature. He found another change overnight to make it even faster, and became the surprise winner of the first race!

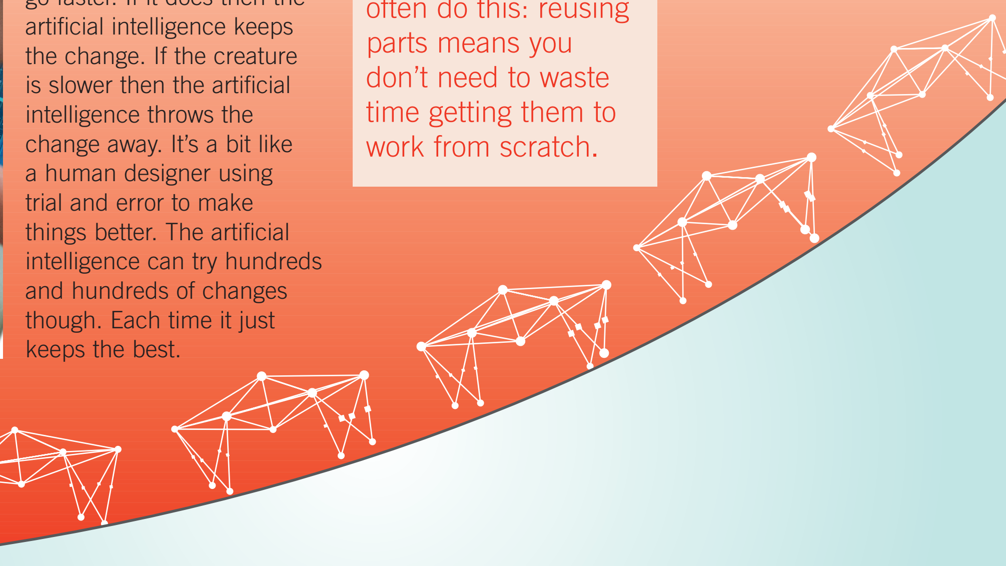

# ai for real

There are lots of ways the approaches to creating artificial intelligences that we've explored in this booklet are being used for real. Here are a few. The programs may not really be intelligent but they can definitely do clever things.

**Expert AIs**
Programs using rules like our noughts and crosses player have helped doctors decide which antibiotics to give a patient. There's also an expert system that helps engineers design airplanes, and another that can help farmers decide what crops to plant, and where and when to plant them. A different one helps in oil exploration and yet another helps people work out how ecological their building projects are.

**Copying the brain**
Neural networks have been trained to help give robots the senses of touch and sight. They're also used to help recognise what people say so that humans can talk directly to computers.

Neural nets are also used by banks where they learn how to predict the money market, and can even automatically sell shares when it looks like their value is about to go down.

Neural nets have even been used to help DJs mix records in clubs. The networks learn the best way to combine the music.

## Trial and Error

Trial and error methods have been used to evolve games for mobile phones. This helps tune them to just the right level of difficulty.

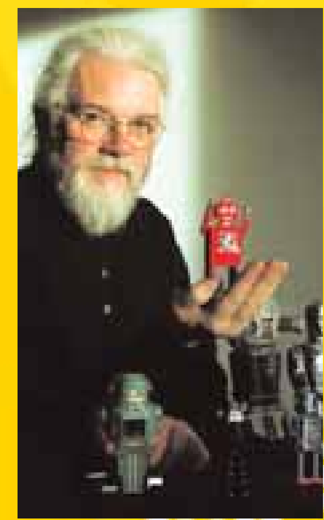

## TV Stars

Robots have even starred in their own hot tv show, Robot Wars, where robots built by teams up and down the country battled for supremecy in the arena under the watchful eye of judges including computer scientist Noel Sharkey.

# more real ais

## See like a fly
Researchers have used some ideas from examining how houseflies see to build silicon chips that can detect how fast things move. The chips are used to help make cars safer by avoiding crashes.

## Driving home alone
In 2005 the American government sponsored a car race through the rugged Mojave desert with a 2 million dollar first prize. The cars in the race were all driven by artificial intelligence software – no humans allowed. The winning robot car, from Stanford University, was called 'Stanley'. It drove itself over the 130-mile desert racetrack in about seven hours.

## AIs making movies
Artificial intelligence software is used in movies, computer games and TV too. When directors want gigantic battle scenes or to show lots of people in a crowd, AI software is used to animate all the computer-generated characters. Each of the characters can be doing something different, and because they have simple artificial intelligence built in, they can react to the characters around them too.

## Helping the disabled
Researchers are looking at ways to connect electronics to the nerve cells that send messages all through your body. In the future this could help people with disabilities link with technology to have a better quality of life.

For example, the silicon cochlea is an electronic device that works in the same way as part of your ear. It connects straight into the brain, and has been used to help deaf people hear.

## The PS3 supercomputer
The Sony PS3 game console allows users to donate some of its spare computing power when it's not being used. In 2007 more than 600,000 PS3 users donated some power this way. Using the Internet to join up all these machines, they created a record breaking worldwide supercomputer to help scientists examine how proteins get their shape. Even with all those PS3s linked together, it's been estimated that the resulting supercomputer was still slightly less powerful, in terms of raw calculating power, than one human brain.

# the end?

It used to be that when we saw artificial intelligences in movies they wanted to take over the world – they were the enemy of humanity. In more recent movies like Wall*E and AI, they are the good guys. Maybe the kinds of AIs we see in movies tells us something about what people at the time think about computers. Being intelligent does not mean wanting to destroy things, and if computers do become more intelligent than us then hopefully they will realise it!

In the meantime they are helping us more and more, making it easier for us to do things, and looking after us. Someday soon they're even likely to become our friends…

This booklet was created for a 2008 Artificial Intelligence workshop series at Chingford Church of England Junior School with support from the Royal Society on a Partnership grant together with funding from Queen Mary, University of London.

It is part of the cs4fn project (www.cs4fn.org) supported by EPSRC, the BCS, HEA and industry. cs4fn is a free website and magazine that aims to inspire students of all ages about computer science. Have fun!

Queen Mary
University of London

EPSRC Engineering and Physical Sciences Research Council

Partnership Grants