

# How Smart is AI?

## Teacher Notes

### The Turing Test

Artificial Intelligence (AI) is capable of generating responses which mimic human responses, but how 'good' are these responses and can we tell if they are computer generated, or from real humans?

The Turing test was developed by Alan Turing in the 1950s to test a machine's ability to exhibit intelligent behaviour and mimic human responses. It takes the form of a list of questions that are asked of two participants. Their answers are analysed to decide on their identity. The idea is that the answers from the AI participant will stand out from a human participant as not being 'quite right'.

Various computer programs have been developed which appear to pass the Turing test, for example ELIZA which uses pattern matching to hold a conversation. You can 'chat' with ELIZA online via various websites. Modern 'chatbots' on the internet are versions of this program.

#### **Activity** - Conversation with a Computer

This is a game where students must try to decide who is a human and who is a 'computer' by asking and analysing a series of questions. The human will be free to answer the questions as they wish, but the 'computer' will need to stick to a list of pre-programmed responses, just like a computer does when they run a program.

Four volunteers (you could use students or staff) need to be briefed before the session on their roles so they know how to behave during the game. There are 4 roles:

- Gina
- George
- Harry
- Connie

The names are important as they will help you remember what roles the people are playing, so if you need to change them please stick to names that start with the same first letter.

Gina and George:

These are the go-betweens. They will take the questions to the participants and bring back the answers.

Harry:

Will give spontaneous human answers to the questions

Connie:

Will give computer answers to the questions from a script.

*Activity from "Computer Science Unplugged" © Bell, Witten, and Fellows, 1998*

Gina and George are there to ensure the test is carried out fairly, they need to relay the questions and answers to Harry and Connie without letting anyone know which is which. Harry and Connie need to be in separate rooms to the 'audience'. It is important they do not identify who they are interacting with, for example by saying "*She said that....*" Harry must give their own short answers. Connie answers by looking them up in the table they are given. Where there are instructions in italics, Connie will need to work out the answer. The questions and answer sheets are given in the Activity Sheets.

## Game instructions

- 1) Before playing the game ask the students if they think that computers are intelligent, or they may be one day. Do they have any ideas on how you would decide whether a computer was intelligent?
- 2) The intelligence test they are about to do is based on asking questions to a human and a computer and seeing if you can tell the difference between them. If the class cannot tell the difference between them then the computer passes the test.  
'Gina' and 'George' are going to take their questions to two people. One will give their own answers, while the other will give answers that a computer might give. Their job is to work out which is which.
- 3) Show them the list of possible questions and ask them to decide which one to ask first. Which one would be good at distinguishing the human from the computer? 'Gina' and 'George' can take the question to 'Harry' and 'Connie' and write their answers down on the Q&A slips (see Activity Sheets).
- 4) When 'Gina' and 'George' come back the class can discuss if they think they can tell who the computer is.
- 5) Repeat the process of asking questions until the class think they have discovered who the computer is.

## Notes

The answers that Connie is reading from are not unlike the ones that some “intelligent” computer programs can generate. Some of the answers are likely to give the computer away quickly. For example, no-one is likely to recite the square root of two to 20 decimal places, and most people (including, perhaps, the children in the class) would not be able to answer that question at all. Some questions will reveal the computer when their answers are combined. For example, the “Do you like. . .” answers sound plausible on their own, but when you encounter more than one it becomes apparent that a simple formula is being used to generate the answers from the questions.

Some of the answers indicate that the question was misinterpreted, although the class might reason that the person could have made the mistake.

Many of the answers are very bland, but safe, and a follow-up question would probably reveal that the computer doesn’t really understand the subject. Answering “I don’t know” is reasonably safe for the computer, and might even make it seem more human—we would expect a child to answer “I don’t know” to some of the questions too, such as the request for the square root of two. However, if a computer gives this answer too often, or for a very simple question, then again it would reveal its identity. Since the goal of the computer is to make the questioners think that they are dealing with a person, some of the answers are deliberately misleading—such as the delayed and incorrect answers to the arithmetic problem. The questions and answers should provide plenty of fuel for discussion.