Before-reading questions

1 ***Example answer***: The idea of AI shown on TV, in films, books and in the media is of machines that are self-aware, conscious and autonomous; that we will not be able to control, and that could be harmful to us.

2 ***Example answer***: You can find examples of AI in everyday life in: software agents like Apple's Siri, Amazon's Alexa, Google's Assistant and Microsoft's Cortana apps; wearable devices and smartphones; machines in hospitals; partially driverless cars; banking systems; photo-management applications like Microsoft's CaptionBot; automated translation programs like Google Translate; social media applications; automated drones; augmented reality applications like Google Glass; and Deepfakes.

3 *Reader's own answer*.

4 *Reader's own answers*.

During-reading questions

CHAPTER ONE

1 Alan Turing was a British mathematician, who first invented the computer, and then the field of AI itself.

2 The Turing test is a test to decide whether a person or a computer program is answering questions typed by a human interviewer.

CHAPTER TWO

1 A modern computer can follow up to 100 billion instructions every second.

2 Humans are much slower and they make mistakes.

CHAPTER THREE

1 An American academic, John McCarthy, was applying for money for a summer school in 1956 and he had to give it a name.

2 Artificial also means fake, which can sound negative; and many of the tasks that AI researchers work on do not require any intelligence to do them.

CHAPTER FOUR

1 SHAKEY could perceive its environment, understand where it was and what was around it. It could receive tasks from users, plan and perform them. It couldn't perceive obstacles without the help of a specially painted and lit environment, and it couldn't work out how to do a task immediately. It couldn't have been used on any practical problem.

2 Search is a basic AI problem-solving technique, which considers all possible plans of action. It is often used for programs that play games like chess, and for satellite navigation systems.

CHAPTER FIVE

1 Heuristic search was first used in AI by IBM worker Arthur Samuel, in the mid-1950s, to write a program to play the game of checkers.

2 An NP-complete problem is a problem for which it is hard to find solutions, but where it is easy to check whether or not you have found a solution.

CHAPTER SIX

1 Knowledge-based AI is human knowledge about a problem, explicitly captured and used within an AI system and usually based on rules.

2 Logic-based AI is important because the problem of building an intelligent system is reduced to writing a clear, logical description of what a robot should do.

CHAPTER SEVEN

1 The two groups were: McCarthy's group at Stanford University, with an old-world view of logic, knowledge representation and reasoning; and Brooks' group at MIT, with a new-world view of AI, based on systems in real-world situations, without explicit knowledge and reasoning, and interacting with their environment.

2 Brooks developed a theory known as behavioural AI and he designed a vacuum cleaner robot to show how it worked.

CHAPTER EIGHT

1 HOMER was a simulated robot which worked in the sea and did tasks like collecting parcels and moving them from one place to another.

2 In 2011, an app called Siri appeared, developed for the Apple iPhone. It was a software agent that users could interact with in natural language, and which could do simple tasks for them.

CHAPTER NINE

1 The new challenge for multi-agent systems was knowing what action to perform for a user and considering what other agents are likely to do when making its decision.

2 In 1997, the company IBM was able to show that an AI system, called DeepBlue, could consistently beat Russian chess champion Garry Kasparov. By the end of the 1990s, programs to solve SAT problems became powerful enough to be used on industrial-size problems.

CHAPTER TEN

1 Machine learning programs are ones which can learn how to do things without being explicitly told how to. A typical application is text recognition.

2 It is the problem of deciding what you need to include in your training data. The more features you include, the more training data you need to give the program. The program will learn more slowly. But if you only include a small number of features in your training data, you may leave out features needed to allow the program to learn correctly.

CHAPTER ELEVEN

1 *Breakout* was one of the first video games to be developed. After playing the game hundreds of times, a machine learning program became expert, learning the best way to win, without anybody telling the program how to do this.

2 The two areas are: image captioning programs because they can correctly identify a key element of a picture and recognize where it is, but cannot really understand the picture; and automated translation because when it tries to translate complex language it translates the text in an unnatural way, using strange sentences which are difficult to understand.

CHAPTER TWELVE

1 ***Example answers***: AlphaFold is a system to help treat serious medical conditions like Alzheimer's disease, using machine learning techniques to predict future signs of the disease.
There are wearable devices, which check how fast your heart is beating or your body temperature, and perhaps in future will be able to identify symptoms of disease or even call an ambulance.
Smartphones are able to recognize and record changes in people's behaviour, which could diagnose dementia long before there are any other signs.
AI techniques can be used to automatically identify eye problems, using neural networks.

2 The main problems for driverless cars are dealing with unexpected events, the fact that they can confuse and frighten human drivers, and they can't predict human behaviour and interact safely with humans. The possible solutions are to use the technology first in areas without many people, on well-mapped city roads, and in driverless car lanes.

CHAPTER THIRTEEN

1 The Singularity is a theoretical point at which computer intelligence goes beyond that of humans and computers use their own intelligence to improve themselves. But software, such as machine learning programs, improves at a much slower speed than the hardware. So even if AI systems did become as intelligent as people, they would not necessarily be able to improve themselves at a speed beyond our ability to understand.

2 The Trolley Problem is an ethics problem about whether to allow a tram, which is out of control, to kill five people or just one person. It can be applied to driverless cars because they will have to make similar life or death decisions, but it is not clear if this is right or fair.

CHAPTER FOURTEEN

1 Office or factory jobs, and jobs involving the movement of people or things, might all be in danger. Jobs involving ideas, requiring strong social skills, perception and dexterity, should all be safe.

2 ***Example answers***: The arguments for the use of autonomous drones in war are that: they could be designed to behave more ethically than human soldiers; it would be better to have robots in wars rather than people; and ordinary weapons are not ethical either.

The arguments against the use of autonomous drones in war are that: they could act largely without human guidance or intervention; they could decide whether or not to take human life; a country with autonomous weapons might also take the decision to go to war more easily and more often, without any risk to its people.

CHAPTER FIFTEEN

1 ***Example answers***: Things that could happen: an AI program that is biased might stop a bank lending money to a certain group because of where they live or the colour of their skin; and because there are more men than women working in AI, systems might be designed for men.

Things that have happened: a Google photo classification system put the label "gorillas" on pictures of black people; and the TIMIT spoken-word data used to train speech understanding programs contains more male than female voices, so they do a worse job of understanding female voices than male ones.

2 AI helps to spread fake news by working out your preferences from what you say you like, the comments that you leave, and the links you follow on social media. It then uses these to find new items that you will also like, and this can change our beliefs either on purpose or accidentally.

CHAPTER SIXTEEN

1 Nagel's test is a philosophical test which considers whether the question "What is it like to be an X?" is meaningful when applied to different things (from humans, to other animals, to objects).

2 The Chinese room scenario is that of a man in a room, who is given written instructions on cards in Chinese, a language he does not understand. It is a kind of Chinese Turing test. The man's intelligence is used only to carefully

follow the instructions he is given, just as a computer would. In the same way a system based on strong AI, capable of passing the Turing test, cannot be produced just by following instructions.

After-reading questions

1 *Reader's own answer*.

2 **a** ***Example answer***: Strong AI is programs that really do understand in the way that people understand; weak AI is programs that show the same ability as humans, but without real understanding.
b *Reader's own answer*.

3 *Reader's own answer*.

4 ***Example answer***: A deep learning program can: identify an illness, but cannot explain its diagnosis; it can decide not to lend money to a customer, but not tell you why; it doesn't always work and can wrongly classify images or data, which could cause mistakes or accidents; it can correctly identify a key element of a picture and recognize where it is, but cannot really understand the picture; it can translate basic text, but not the complex language of a novel; it can't understand our preferences, if we do not understand them ourselves.

5 *Reader's own answer*.

6 *Reader's own answers*.

7 *Reader's own answers*.

Exercises

CHAPTER ONE

1 **1** David Hilbert **2** Alan Turing
3 Konrad Zuse **4** An American team
5 John von Neumann **6** Joseph Weizenbaum

CHAPTER TWO

2 **1** c **2** f **3** b **4** e **5** d **6** a

CHAPTERS TWO AND THREE

3

AI systems	Computer games
SHAKEY STRIPS	*Call of Duty* *Minecraft*

Computer programs	Computer programming languages
PowerPoint Word	Java Python

CHAPTERS FOUR TO SIX

4 **1** If we can build programs to solve problems that people find hard, then surely this **would** be an important step on the road to AI.
2 If there is a solution, we **will find** it in the end using this process.
3 If you **did** a quick experiment, you would see that in almost all possible combinations of the puzzle there would be three ways to move.
4 Now suppose you **have** 100 people, and you need a team of fifty. Then you would have to check 100 billion billion billion possible teams.
5 This means that if you could find a quick recipe, or program, for solving just one NP-complete problem, then you **would have found** a recipe for solving all of them.
6 It simply depends on whether you **can accept** the pattern of reasoning, and the conclusions you come to, if the premises were true.

CHAPTER SEVEN

5 **1** b It avoids obstacles if it perceives them.
2 d It shuts down if it has a low battery.
3 e It empties the dirt container if it is full.
4 c It returns to dock if the battery is low or the dirt container is full.
5 f It switches on the vacuum when it perceives dirt.
6 a It chooses a direction randomly and walks/moves in that direction.

CHAPTER EIGHT

6
1 Yes.
2 I'll go to the port in order to pick up the parcel.
3 8.56 p.m.
4 On the boat.
5 No.

CHAPTER NINE AND TEN

7 1 f 2 d 3 e 4 b 5 c 6 a

CHAPTER ELEVEN AND TWELVE

8
1 true 2 false
3 true 4 true
5 false 6 false

CHAPTER THIRTEEN

9
1 **The Three Laws of Robotics** were created in a series of stories about robots with a kind of strong AI.
2 **Consequentialist reasoning** decides on the ethics of actions based on their results.
3 **The greatest happiness principle** says that one should act for the greatest good for society.
4 **A general rule of good actions** means not taking any action, because taking any life is wrong, even though not acting leads to five other deaths.
5 **Virtue ethics** identifies a decision-maker with good ethics and tries to act like they would in the situation.
6 **The Asilomar principles** was one of the first and most important systems of ethical AI, which had twenty-three rules.

CHAPTERS FOURTEEN TO SIXTEEN

10
1 As a general-**purpose** technology, AI can be misused or have unintended consequences.
2 Since 2001, remote-**controlled** drones, flown from a distance, have been used by the USA in wars in countries like Afghanistan, Pakistan and Yemen.
3 As machine learning systems have been applied to more areas, we have begun to understand how automated decision-**making** systems can show algorithmic bias.
4 Sometimes, the bias is explicit, such as in the TIMIT spoken-**word** data widely used to train speech understanding programs.
5 These types of applications are called augmented reality: they take the real world and cover it with computer-**generated** information or images.
6 Strong AI, building machines that really have conscious minds, self-**awareness** and understanding, like us, is still a distant goal.

Project work

Reader's own answers.

Essay questions

Reader's own answers.