1

WHAT IS AI AND WHY AI MIGHT BE USEFUL IN EDUCATION?

How can AI benefit education? What is AI and how can I use it effectively? What sort of AI do I need? These might be some of the questions you are asking yourself. Before we dive into these questions, let's take a moment to get clear on *what* AI is and *why* AI might be useful. We will also try to enthuse you about AI's potential for use in education.

WHAT IS AI?

The term "Artificial Intelligence", abbreviated to AI, is not something for which there is any single accepted definition. Many scientists disagree about exactly what the precise definition of AI should include. As this chapter will explain, there are two main types of AI: machine learning and Good Old-Fashioned AI (GOFAI). Some people believe that only machine learning should be called AI, but many others believe that the definition of AI should also include tools and technologies that make intelligent decisions in other ways. We have selected a simple definition of AI in an attempt to include the vast array of different sorts of tools and technologies that could be considered to come within the bounds of the phrase "Artificial

DOI: 10.1201/9781003193173-1

Intelligence". The definition we have selected is taken from the *Oxford English Dictionary*:[1]

> The capacity of computers or other machines to exhibit or simulate intelligent behaviour.

A SHORT HISTORY OF AI

MACHINES THAT COULD BEHAVE, BUT NOT INTELLIGENTLY

With a definition in place, let's look back at where AI has come from. For many centuries, humans have been intrigued by the task of creating representations of living creatures, including humans. These representations are often referred to as automata and they date back to the Middle Ages, possibly even earlier. In the 19th and early 20th centuries, automata reached the height of their popularity. From bears that could turn somersaults to magicians who could see another automata in half, and nightingales that sang in golden cages, these party pieces were increasingly intricate and could perhaps be considered forerunners of AI.

Or perhaps more accurately, we might consider them to be the forerunners of the field of cybernetics, which is a scientific area of study that explores the control and communication that happens in animals and in machines. The study of cybernetics was started by Norbert Viner in the mid-20th century and is still very much at the heart of what robots can do. These cybernetic robot predecessors were more mechanical than intelligent, but their creation influenced the field of robotics that evolved. Even today, not all robots are intelligent; some are just labour saving through their speedy completion of mechanical, repetitive tasks. However, many robots are also intelligent and are part of the human desire, played out over time, to create objects that can behave in intelligent ways.

We also love to tell stories about objects that behave intelligently, and robots have long been a love of the film-making industry. Who cannot but be endeared by robot characters like Star Wars' C-3PO,[2] or WALL-E,[3] feel fearful of Blade Runner's Roy Batty[4] or The Borg,[5] and simply be amazed by Ava in Ex Machina![6] The reality is nothing so dramatic. It is certainly true that intelligent software and robots are a reality, but none of them have the all-round capabilities of their movie-star peers. The ability of AI systems to achieve more than one area of expertise is still mainly a fantasy. From tangible robots to invisible software, AI is a specialist operator with no ability to transition from one area of expertise to another. A self-driving car cannot play chess; a surgical robot cannot drive a car. It's worth noting, however, that a human surgeon can likely drive a car and play chess and a great deal more besides.

But let us return to our brief look back at the history of AI.

Another important person in the history of AI is Alan Turing who, in 1950, wrote a famous article titled "Computing machinery and intelligence" in which he posed the question: "Can machines think?"[7] Alan Turing was a mathematician and code breaker at Bletchley Park during the Second World War. Turing proposed a clever test that could be used to decide if a machine was thinking and, therefore, was intelligent. This test, which is called the Turing Test, presents the proposition that if a computer can fool a human into believing that it is really a human, then that machine deserves to be called intelligent. This thought experiment captured the interest of many scientists and helped to progress the birth of AI. Indeed, there are still Turing Test challenges today in which computer scientists pit their AI against one another to see whose system can convince the most people.

THE TEN MEN WHO GAVE BIRTH TO MODERN AI

Following the publication of Turing's famous article, the field of AI evolved at a rapid pace, and in 1956 a momentous meeting took place at Dartmouth College in New Hampshire, in the United States.[8] A ten-man group of scientists met with the aim of studying human intelligence in all its richness and from all aspects. Their goal was to be able to describe each feature of human intelligence so precisely that a machine could be built to simulate it. The scientists believed that they would be able to make significant advances towards their goal over the period of a summer.

They soon discovered that human intelligence was *far* more complicated than they had understood and progress during that summer was small. Nonetheless, the occasion of this meeting was extraordinarily significant, as it gave birth to what we now recognise as the scientific discipline of AI.

It is interesting to wonder had the ten scientists been from more diverse backgrounds, if their group would have taken AI in a different direction. We will never know. What we can be sure about is that the lack of diversity amongst those who work in AI is an ever-present concern. We hope that by making AI accessible to more people, a more diverse population might become interested in working with AI.

From that date onwards, the task of creating computer programs that behaved in intelligent ways was the cutting edge of science. At this stage, AI was not focused on robotics, but on developing software that could enable computers to interact intelligently. Early attempts were very simple. Systems such as ELIZA,[9] a computer program that played the role of a psychotherapist, were text-based and required the person playing the role of the patient to type out their problems and questions. The ELIZA software was programmed to look for keywords in what the patient typed. When a keyword or phrase was found, the software triggered a stock answer template and ELIZA offered her advice in the text on the computer screen. It is hard to believe that anything so crude could fool anyone into consulting with ELIZA for more than an initial sentence. But several people were duped by ELIZA, at least for a while.

GOOD OLD-FASHIONED AI AND EXPERT SYSTEMS

If those ten scientists who met in New Hampshire are considered the fathers of AI, then maybe ELIZA should be the mother.[10] The important thing about systems like ELIZA is that it identified a particular approach to simulating intelligent behaviour. This approach is called production rule-based pattern matching and ELIZA "gave birth" to many similar systems over the decades following her inception in 1964. In fact, the production rule-based systems that evolved from ELIZA became sophisticated enough to accomplish advanced activities such as diagnosing an illness from a set of symptoms and suggesting the treatment regime based on these symptoms. These systems are referred to as expert systems and were used in a variety of different fields, such as medicine.

The pinnacle of the GOFAI movement came in 1997 when a system built by IBM, called Deep Blue,[11] beat the then chess grandmaster Gary Kasparov at the game of chess. This was extremely impressive and marked the high point of this phase in AI's history.

The problem with these Good Old-Fashioned AI (GOFAI) systems was that the actions the AI was able to perform had to be pre-programmed into the software when it was written – a huge task. For games like chess, there may only be a certain number of moves that each particular chess piece can take, but there are many millions of iterations that the combination of these moving chess pieces can create. In fact, trying to look ahead for just two moves in a chess game would generate 1,225 possible chessboard states. Looking ahead 20 moves will generate 2.7 quadrillion possible board states.

To write a computer program to find ways of dealing with the existence of all these possibilities and to be the best in the world at it too is no mean feat. However, there is a severe limit to the intelligence that this style of AI could achieve. Once the knowledge was written into the computer program code, the system could not be updated without going back and changing the code. No matter how many disease cases they diagnosed, or gas pipe fractures they identified, or games of chess they played, GOFAI systems will never improve.

Before we confine GOFAI to the history books, however, it is worth thinking about the uses it still affords. As teachers, we can easily see how a GOFAI system could still be extremely useful in the classroom. For example, planning a school trip involves many possible steps, as illustrated in Figure 1.1. An easy-to-use app that helped

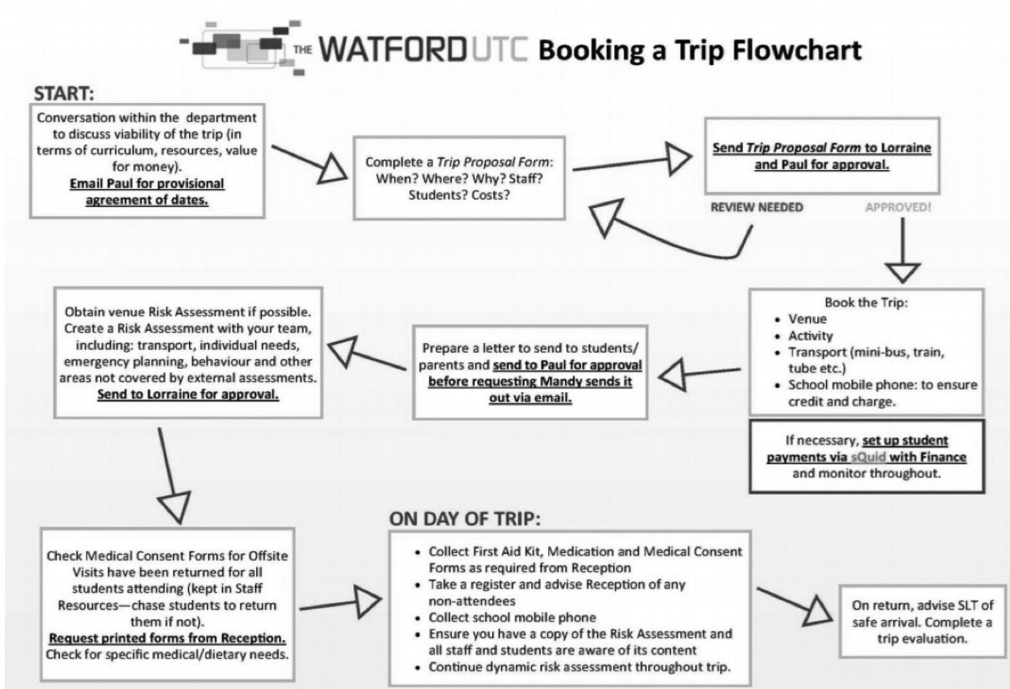

Figure 1.1 An example of the steps involved in planning a trip.

teachers step through all these processes and decision points could be extremely helpful and it could be built using GOFAI.

It is also worth taking a moment to note that there were suggestions that production rules of the sort used to create systems such as ELIZA and her "offspring" could represent the basis of human thought. For example, the work of John Anderson who developed the adaptive control of thought (ACT) theory of human thought as a set of production rules.[12] As we judge the wisdom of his work, we should remember that we know a great deal more about human thought now than we did at the time that John Anderson was developing his ACT theory. We should also remember that his work precipitated a generation of cognitive STEM tutoring software that was extremely successful at tutoring students in maths and science. Indeed, these cognitive tutoring systems are the basis for the systems that are sold today by the spin-off company from Carnegie Mellon University called Carnegie Learning.[13]

In truth, the Carnegie Learning example was not the zeitgeist in the late 20th and early 21st centuries. From the zenith that was Deep Blue came several decades where AI moved at a much slower pace than it had in the years immediately following the Dartmouth College meeting. The technical limitations of machines that could not learn severely restricted what could be achieved with AI, and funding was less available. The deepest of the AI winters to date was upon us.

MACHINES THAT CAN LEARN

It is always darkest before the dawn so they say, and the AI winter was indeed quite a gloomy time to be working in AI. Progress in AI was slow as we said farewell to the 20th century, and the new millennium dawned. Despite the coolness towards AI, the drive to produce AI systems that, like humans, could become better and better and better at a particular activity was a strong motivator for many aspiring computer scientists. In 2011, Google formed Google Brain and thus emerged the field of machine learning and started to grow.[14]

Machine learning is the approach that many modern AI systems use. Indeed, most of the AI that we use daily is likely, at least in part, to be using machine learning to produce the behaviours that make the system seem intelligent. Imagine, for example, that you have been to a social gathering and seen one of your friends wearing a particular pair of shoes that you really like. Your friend is being coy about where these shoes came from, and how much they cost. But you really want to find out so that you, too, can purchase these shoes whilst at the event. You manage to persuade your friend to let you take a picture of the shoes. This means that you can use that photograph with AI software that is capable of visual search. This type of AI system will search millions and millions and millions of images to try to find an image that matches the shoes in the photograph.

With a bit of luck, the shoes you desire are available on e-commerce sites. If so, it is highly likely that the visual search AI system will be able to find them for you.

Visual search is not just for shoes. It can be a real boon for a teacher too. It is a certainty that, as a teacher, you will be caught out by a question, often a simple one, that you cannot answer. Taking children out to explore, collect, and identify different species of plants, flowers, or trees before sorting and classifying is a common activity. In the vast majority of cases, a child will ask a teacher to name a species of flower, for example, that they may be unfamiliar with. However, help is on hand as there are many apps that utilise visual search AI so that the teacher can take multiple photos, upload them, and get an instant identification. As we said at the start, making friends with AI is a good idea.

TRAINING THE SYSTEM

Systems, such as the visual search systems just described, use machine learning. The instructions within a machine learning AI system, sometimes referred to as the algorithm, need to be trained.[15] In the same way that we as humans can be trained to recognise similarities and differences, AI algorithms can be trained.

The training that enables humans to recognise particular pairs of shoes may not be explicit but it happens as we grow and become more able to see and process the world around us. Every day, as a child grows, they see millions and millions and millions of images from the world around them. The child gradually becomes more accurate at recognising things, mum's face, their own hands, the picture of a whale on the wall next to their cot, for example. In a similar way, the machine learning algorithm must be trained using millions and millions of images of pairs of shoes, or a plant, for example, in order to enable that algorithm to accurately recognise a particular pair of shoes or a plant, when an image is input to the system.

Machine learning AI systems can do many things, not just image recognition visual search. But visual search is one of the activities that machine learning AI is becoming extremely accurate in. Indeed, many of us helped to train these algorithms, mostly, without realizing that we were doing so.

Next time you are signing into a website, and you are asked to select all the instances of a particular item from a set of images (such as staircases, or bridges, or traffic lights),[16] you are helping to label a photograph as containing the item that an algorithm is being trained to recognise. And notice that traffic lights, bridges, staircases, etc. come in a variety of shapes and sizes. Training algorithms to recognise what an image is illustrating is not an easy task. As we will show you in Chapter 5, a great deal of time, and often a great deal of human effort too, is required to prepare the data before a machine learning algorithm can be trained.

PREPARATION IS ESSENTIAL AND HUMAN TIME INTENSIVE

The preparation phase that happens before the machine learning algorithm can start to learn from data requires that data are labelled correctly. Training in order to learn is not just for machines of course. Humans need to do this too. Let's look, for example, at how we learn about shapes. Our world is dominated by shape and space, in the kitchen, in the bathroom, and outside the home. Look around you! It's not hard to see then why geometry is a key part of the component in the programmes of study for maths. Learning about shapes and space enables learners to sort out their visual data in order to think about and describe their environment.

AU: Should it read 'It's not hard to see then why geometry is a key component in the programmes of study for maths.'

To begin with, children need to have the language, or labels, to name and describe shapes. This normally begins through play. Naming the blocks, spotting and matching the shapes, shape hunts around the school and the local environment are all necessary precursors to understanding future work on measurements, construction perspective in art, and so much more.

This layering of concepts is replicated in machine learning. If the data from which the machine learning algorithm is going to learn are in the form of images and the machine learning needs to be able to identify which of these images contain traffic lights, staircases, shoes, or whatever, then all the images used to train the machine learning algorithm will need to be labelled according to whether or not they contain traffic lights, staircases, or shoes. Who is going to do this labelling? It may be people who work for the company making the AI product, but it may also be you, or me, or our friends as we have just illustrated. The point to note here is that machine learning needs a lot of help from people to even start to learn anything.

LEARNING IS FOR LIFE, EVEN FOR MACHINES

As we have seen, machine learning Artificial Intelligence is built on learning – learning from examples, or more accurately learning from data. The machine learning AI algorithm is given millions and millions of data examples from which it learns. That is, it learns from experience.

The real power of machine learning is found following the training phase. Once the machine learning system has been trained to a level that its developers are confident it is accurate enough to be used in the real world, the system *continues to learn and improve.*

Again, the same happens with children's learning. Once we have taught them the language of shape and space and they can discriminate and understand the properties of them, they are then ready to go beyond the maths lesson and apply their knowledge to their construction lessons in design and technology. In this way, learners can build upon and consolidate their learning.

TRANSPARENCY AND AI, OR UNDERSTANDING WHAT'S HAPPENING IN THE BLACK BOX

You might be thinking that the advent of machine learning has made GOFAI obsolete, because why use an AI that cannot learn, when you have an AI that can learn? In fact, why even call an AI that cannot learn an AI? There are many people who do not regard GOFAI as AI, and yet it really does come with some significant advantages, despite the rather large disadvantage of not being able to learn. GOFAI creates a trail of decision-making points as its rules are fired. Therefore, any decision that a GOFAI system makes can be explained. This makes GOFAI systems highly transparent.[17]

Machine learning AI, on the other hand, has no rules. It can therefore be extremely hard to know precisely why a machine learning system has made any particular decision. Machine learning AI systems are what we call black box systems.[18] There is no transparency available to show us what the algorithm in the machine learning AI system has done or why it has reached a particular decision. Machine learning is fast, and it learns. It has many advantages over GOFAI. However, the disadvantage of not being able to provide an explanation represents a significant problem, particularly for education and training activities, where it is really important to be able to explain, and often justify, why a particular decision has been made.

In educational settings, it would be a total disaster if you could not explain or justify actions taken. Imagine a situation where peer-on-peer abuse has taken place. The school has investigated and decided upon a course of action that affects one child more than another. In this case, it's your child. On arriving home, the child paints a picture which in your opinion does not justify the action the school has taken. Against this background, and with limited information, every parent would expect and deserve to understand the justification for the decisions made if we expect them to support the actions taken.

GOFAI AND MACHINE LEARNING CAN WORK TOGETHER

The machine learning community is keen to rectify the lack of transparency associated with machine learning AI. Across the world, AI scientists are working on what is often referred to as Explainable AI (XAI).[19] This goal of XAI is to increase the transparency of machine learning AI and find ways to generate explanations and justifications for the decisions these systems make. This may mean blending GOFAI and machine learning techniques to try and get the best of both worlds.

One of the other challenges that machine learning AI faces is the need for enormous amounts of data from which the AI can learn. As an example, let's take an essay grading machine learning AI system. To grade essays accurately that machine learning AI would need to have processed, or "experienced", millions of essays across the full range of possible grades before it can grade essays accurately. Do we have millions of examples of graded essays that cover the full range of possible marks? Are they in a digital format that can be labelled and made accessible to the machine learning AI? We may be able to collate sufficient examples, but it is not easy. The hefty data requirements of machine learning AIs are a key restriction on their application. Not surprisingly, AI scientists are keen to find a way to address this challenge and a new field of machine learning called TinyML[20] is evolving. It is early days for TinyML and it will be interesting to see how successful this new technique may be.

A key component of engaging with AI is to understand the importance of data to machine learning AI. Think about the data you have access to and make a list.

List three kinds of data you have access to.

And

List three tasks you might like an AI to help you with.

As you read this book, think about ways in which an AI might usefully process your data. In the next chapter, we will start to explore the sorts of tasks that AI is better at than we humans are and the tasks that we humans are better at than AI. It would also be useful for you to start thinking about the sorts of tasks that AI might be able to take on to help you.

The data that you think about could be data about individuals, e.g., how they are performing academically. Or perhaps it is data about how those individuals are feeling, whether they are anxious, or whether they are feeling supremely confident. There are many sorts of data available in any organization, and quite often, some of the most obvious kinds of data are not really thought about when we ask the question, what data is available? For example, data about the temperature of a classroom, or about the time it takes people to get from one part of the building to another, or data about light levels in each classroom, or about which pupil regularly sits next to which other pupil.

We appreciate that you are time-poor and thinking about data sources may not seem like a good use of your time. Trust us, it is. In particular, if you are a school leader with many competing demands on your time, it is easy to overlook the value of data. We know that the very fabric of the building and all its resources are important assets in supporting the delivery of learning. To this end, schools have inventories, logbooks, and financial data that document the purchase of everything from PE equipment and technology to the replacement of carpets and even the number of chairs that are ordered. Schools are a treasure trove of underutilised data. Dealing with issues as they arise can be costly and other budgets can suffer as money is pulled from one source to support another. Therefore, schools need to really examine all the data they have, as often the solution to many of the logistical nightmares schools have can be found in the data they hold. And once you know more about your data, you will be much better able to see how AI might help you, through an intelligent repair and maintenance system perhaps. It really is always worth thinking about the sort of data that is and could be available to you in order to make good decisions about where best to deploy AI.

AUTONOMY AND ADAPTIVITY

In this chapter, we have presented a very brief description of the historical roots of today's AI. We have described two main types of AI: GOFAI and machine learning. One type of AI can learn; the other cannot. One type of AI can easily generate explanations for all decisions made, the other cannot. The key concept of data has been delineated, which provides machine learning AI with the experience from which it learns. For the machine learning AI to learn from this data, it must be prepared by people, an activity which many of us participate in unwittingly every day. Finally, in this chapter we want to prepare you for the chapters that follow by introducing you to two characteristics of AI, both machine learning and GOFAI (to an extent), that you will encounter as you read on. These two characteristics are *autonomy* and *adaptivity*[21] and you will see them referred to repeatedly in the rest of this book.

Autonomy is what allows AI systems to complete actions without constant guidance from humans.

Of course, it is important to bear in mind that on some occasions, even if an AI system is capable of behaving autonomously, it may not be desirable for it to do so, for example, autonomous weapons. A balance between autonomy and human guidance is often the best way forward.

Adaptivity describes the way that AI can interact with a person, perhaps to help them learn arithmetic. Adaptivity changes the way that the AI interacts based on the actions the person takes.

In the example of a pupil learning arithmetic, the AI might adapt the difficulty of the arithmetic activities that the pupil is asked to complete, and it might also offer less or more help and support depending upon how easily the pupil is completing the activities. As you read, and the sorts of data that you might have, think about autonomy, and adaptivity too. How could an AI that does not need constant guidance and that can adapt its behaviour be helpful to you?

USEFUL RESOURCES

- Anthony Seldon and Oladimeji Abidoye. *The fourth education revolution: will artificial intelligence liberate or infantilise humanity*. The University of Buckingham Press.
- Rosemary Luckin. *Machine learning and human intelligence: The future of education for the 21st century*. UCL Institute of Educational Press.
- Rose Luckin and Wayne Holmes. *Intelligence unleashed: An argument for AI in education*. Open Ideas at Pearson.
- Barbara Means, Robert Murphy, and Linda Shear. *Understand, implement & evaluate*. Open Ideas at Pearson.

NOTES

1 https://www.oed.com/viewdictionaryentry/Entry/271625
2 Star Wars' C-3PO.
3 https://en.wikipedia.org/wiki/WALL-E
4 https://en.wikipedia.org/wiki/List_of_Blade_Runner_characters#Roy_Batty
5 https://en.wikipedia.org/wiki/Borg
6 https://en.wikipedia.org/wiki/Ex_Machina_(film)
7 https://academic.oup.com/mind/article/LIX/236/433/986238 https://scholar.google.com/citations?view_op=list_works&hl=en&hl=en&user=VWCHlwkAAAAJ
8 https://en.wikipedia.org/wiki/Dartmouth_workshop
9 Weizenbaum, J. (1966). ELIZA—a computer program for the study of natural language communication between man and machine. *Communications of the ACM*, 9(1), 36–45.

10 Of course, we jest, and there were many women who created the conditions for, and pushed the boundaries of, modern computing. See, for example: https://blog.re-work.co/female-data-science-pioneers-you-may-not-have-heard-of/

11 Deep Blue also used another popular GOFAI technique called *Search*. There is insufficient space in this book to give thorough explanation, but as the name suggests Search involves working through possible chessboard states in order to decide which to select. The enormity of the search space for the game of chess means that the team who built the AI had to come up with different techniques to reduce the size of the search space that Deep Blue needed to navigate. You can read about the way Deep Blue worked in this article: https://core.ac.uk/download/pdf/82416379.pdf

12 https://en.wikipedia.org/wiki/ACT-R and Anderson, J. R. (1996). *The Architecture of Cognition* (1st ed.). Psychology Press. https://doi.org/10.4324/9781315799438

13 https://www.actuaries.digital/2018/09/05/history-of-ai-winters/

14 https://www.nytimes.com/2016/12/14/magazine/the-great-ai-awakening.html AND https://en.wikipedia.org/wiki/Google_Brain

15 This description is specific to a very commonly used form of machine learning called supervised machine learning We will discuss this more in Chapter 5, along with some other types of machine learning.

16 https://www.techradar.com/uk/news/captcha-if-you-can-how-youve-been-training-ai-for-years-without-realising-it AND https://towardsdatascience.com/are-you-unwittingly-helping-to-train-googles-ai-models-f318dea53aee AND https://www.google.com/recaptcha/about/

17 https://en.wikipedia.org/wiki/Algorithmic_transparency

18 https://en.wikipedia.org/wiki/Black_box https://royalsociety.org/-/media/policy/projects/explainable-ai/AI-and-interpretability-policy-briefing.pdf AND https://en.wikipedia.org/wiki/Explainable_artificial_intelligence

19 XAI.

20 https://www.arm.com/blogs/blueprint/tinyml AND https://www.oreilly.com/library/view/tinyml/9781492052036/

21 https://course.elementsofai.com/1/1