Artificial Intelligence and Machine Learning for Business

Fourth Edition

Steven Finlay

A No-Nonsense Guide to Data Driven Technologies

Artificial Intelligence and Machine Learning for Business

A No-Nonsense Guide to Data Driven Technologies

Fourth Edition

(Sample)

Steven Finlay

Relativistic

e-mail: AI@relativistic.co.uk

No plants or animals were mistreated in the writing of this book.

Cover image: "Mind bots." Thanks to Shutterstock.

Contents

Foreword

Artificial Intelligence (AI) is proving to be the most transformative and disruptive technology of our age. AI-based processes, tools and systems are impacting us as individuals, in business and as a society every day, in almost everything we do.

In many organizations, AI has already had a direct impact on many business activities. This is via automation, more accurate decision-making and enhanced customer service to name but a few. It seems that almost everywhere AI-based systems are replacing or enhancing tasks that were once undertaken by people. This is giving organizations that fully embrace these technologies a competitive advantage over their rivals because of the efficiency savings and improved customer experiences that such systems can help to deliver.

AI may seem to be everywhere and is sometimes marketed as a universal solution to every business problem. However, as with any new technology, it's better suited to some problems than others. Success is not guaranteed. Consequently, it pays to have a grasp of how the technology works, why it's useful and where it can be applied most productively. Most importantly, you need a clear vision as to how the technology is going to improve what you do and deliver products and services more effectively. Blindly adopting a technology without clarity of purpose, and having a clear understanding of the pitfalls and limitations of that technology, nearly always leads to failure. That great lesson from the gold rushes

of history is that it's the sellers of picks and shovels that profit most, not the seekers of gold themselves. To put it another way, many organizations have paid a lot of money to buy expensive AI-based tech, only to then struggle to reap any practical benefits from it.

The goal of this book is to peel away the hype that's so often associated with AI and machine learning. It describes, in simple terms, how AI and machine learning work, and how they can be used to improve services that organizations provide. This is while avoiding some of the issues that often bedevil AI projects such as biased decision-making and data protection issues.

When I wrote the first edition of this book, it was intended as a "Quick read" that could be digested in an hour or two. With each new edition the length of the book has grown as new material and additional chapters have been added and now, this latest edition, is probably around twice as long. However, I have striven to retain the concise no-nonsense style of the original, dispensing with trivia and useless anecdotes that won't add much value because this was a key feature that readers said they liked about it. Therefore, the book may now take a little longer than two hours to read but I hope you find the time well spent.

1. Introduction

The use of artificial intelligence-based technologies has exploded in recent years. AI is being used in almost every walk of life to improve processes and enhance peoples' everyday experiences via "Artificially intelligent" machines and computer interfaces. Amazon's Alexa and Google Translate are two well-known software products that demonstrate the benefits that these technologies can deliver. Likewise, facial recognition systems, predictive texting and setting the interest rates on credit cards are further examples of where AI-based technologies are routinely being applied in the real world.

Many products and services are also adaptive. They tailor their responses to the behavior of individual users. TV and music streaming services learn to identify the content you like and present you with recommendations that you'll no doubt be interested in. Change the type of music you listen to and their recommendations will change too. Likewise, you can buy heating systems that learn to anticipate when it's the best time to turn the heating on so that you don't have to bother, while at the same time optimizing energy usage to reduce your bills. These types of service, that continually learn and adapt, are further examples of artificial intelligence in action.

This concise text provides a managerial (i.e. non-technical and no complex equations) overview of artificial intelligence and machine learning, what they are and how they're used. No prior knowledge is assumed. To put it another way, if you can read and write and do

basic arithmetic (there is a bit of arithmetic, but not that much) then you should be OK with the material in this book.

A good question to ask at this point is: Why do I need to know about these things? One reason is personal. AI has become the primary tool that organizations use to leverage the data they hold about you. They use AI-driven tools to predict how you are likely to behave under different circumstances, and hence, the way they should treat you in order to maximize their (and sometimes your) objectives. They use these systems to decide if you'll receive a great offer or a poor one, if you should be placed at the front or the back of the queue, if you'll be treated as a suspect in a criminal case or how much you can expect to receive when you make an insurance claim. Therefore, it's not a bad idea to know something about these things so that you can understand why an organization may have treated you one way and not another.

The other reason to learn about artificial intelligence, and the one that's the main focus of this book, is that it's now a mainstream business tool. Not that long ago, artificial intelligence was the domain of a few nerdy specialists working mainly in academia or tech start-ups in Silicon Valley. These days, regardless of what business you are in, applications of artificial intelligence can be found across the full range of business activities. This covers everything from short-listing CVs to help HR professionals decide who to hire/fire, chatbots answering customer queries, to robots on the production line, warehouse management and customer deliveries. As a consequence, artificial intelligence is supporting or supplanting humans in many domains.

Artificial intelligence has arrived big time. It's no fad and it's here to stay. Those organizations that are using it to solve business problems, improve efficiency and cut costs are benefiting at the expense of their rivals.

This doesn't mean you need to learn all the things that a technical specialist (a data scientist) needs to know. However, having a working knowledge of what artificial intelligence and machine

learning are, and knowing how they can help your organization to deliver better products and services, will be beneficial. Not least, because in order to make effective use of these tools they need to be focused on business objectives to address specific problems that organizations face.

If, on the other hand, you happen to be an equation quoting, formula juggling, bad ass mathematical genius who thinks they know all there is to know about artificial intelligence and machine learning, then this book may also have value for you too. Possibly, a lot more than you might think. Why? Because if all you focus on are the computational aspects of the subject, then you face a real risk of hitting a brick wall when it comes to delivering useful solutions in the minefield that is the real world; a world populated with social, ethical and political issues. This, together with a growing raft of privacy and data protection legislation, could derail your solutions no matter how good they are mathematically. Without consideration of these "Soft issues" the best case is that the solutions you develop don't get to be deployed. The worst-case scenario is that you design an artificial intelligence-based system that lands you in court because it unfairly discriminates against minorities, women or some other group of people.

Almost everyone acknowledges that artificial intelligence has the potential to deliver immense value for businesses and, almost daily, there's another news story about some wonderful new application of AI that's going to save lives or boost profits. However, the rather sobering finding, from a number of studies, is that:

Most business-focused AI projects fail to deliver[1,2].

The reason why so many organizations fail with their AI projects is predominately due to a lack of understanding of these broader social

and organizational issues rather than any failings in the underlying AI-based technology being deployed.

To get the most out of artificial intelligence, data scientists need to engage with business users to understand their problems. Data scientists also need to understand an organization's culture and its approach to the adoption of new ideas, technologies and working practices. Legal and regulatory issues in the region(s) in which their clients operate also require due consideration. It doesn't matter how good a solution is in terms of cutting-edge hardware and software, if it's not aligned with an organization's business objectives and operational processes, then it's all just a waste of time and money.

Lots of solution suppliers can bamboozle you with their fancy tech and the latest terminology, which is sometimes just a rebranding of last year's tech with a new twist. However, the suppliers who add value are those willing to spend time understanding how you and your organization work. They'll then determine if, and how, their solutions can be used to improve what you do and explain this to you in simple language which you can understand without needing to reach for Wikipedia. A competent AI expert will also be willing to be honest and walk away, if they determine that what you need is something else to solve your problems. AI isn't a universal panacea for all the world's ills and, very often, other tools will be far more suitable for the task in hand.

Successful artificial intelligence is a two-way thing. Data scientists need to know something about your organization and what it does, and you need to understand a little bit about artificial intelligence and machine learning. Without this joint understanding, it's unlikely that you or your organization can realize the full benefits that artificial intelligence has to offer.

OK. So, what will you learn from reading this book? The key topics that we are going to cover in the following chapters are:

- What machine learning and artificial intelligence are.

- The sort of things organizations use artificial intelligence for.

- What predictive models, which underpin almost all real-world AI-based applications, look like.

- The relationship between artificial intelligence, machine learning and Big Data.

- How advanced forms of machine learning are applied to drive artificial intelligence applications such as object recognition and language translation.

- The people, tools and other resources needed to apply artificial intelligence in practice.

- How to use artificial intelligence to improve business processes and the bottom line.

- Legal and ethical issues that need consideration when developing AI-based solutions that are going to be used to make decisions about how people are treated.

- The current limitations of machine learning and artificial intelligence.

A list of recommended further reading is provided in Appendix B, and if you see any text in **_bold italics_**, then that means there is a corresponding entry in the Glossary (Appendix C).

2. What is Artificial Intelligence?

Nearly all real-world applications of artificial intelligence are based on **machine learning**. This is the use of mathematical procedures (**algorithms**) to analyze data with the aim of discovering useful patterns (relationships or correlations) between different data items. Once the relationships have been identified, these can be used to make inferences about the behavior of new cases when they present themselves. In essence, this is analogous to the way people learn. We observe what goes on around us and draw conclusions from our experiences about how the world works. We then apply what we've learnt to help us deal with new situations that we find ourselves in. The more we experience and learn, the better our ability to make decisions becomes.

One, widely used, application of machine learning is object recognition. The goal is to develop a system that can identify everyday objects from images the system is presented with. The data used to develop an object recognition system consists of pictures of different objects such as chairs, umbrellas, washing machines and so on. Each picture presented to the machine learning algorithm is labeled to identify what type of object it contains. For each type of object there may be hundreds or thousands of different images, representing alternative forms of that object from different perspectives (you might be surprised at just how many variants of an umbrella there are).

By analyzing the different images, the algorithms recognize that

certain objects are associated with certain features (patterns). Chairs tend to have protuberances (legs) coming from a flat, often squarish, base. They are also differentiated from stools by having a back rest. Washing machines tend to be cube shaped with knobs on and are almost never pink or orange (with the odd exception). Similarly, umbrellas are usually long and thin (when closed), are often, but not always, black and so on.

The most common application of machine learning is prediction. It's about using machine learning to determine something you don't currently know, based on the information that you currently have available. The patterns that one finds relate to the relationships between behaviors and outcomes. Very often this relates to peoples' past behavior and what they subsequently went on to do. Having identified the relationships that exist, it's then possible to make predictions about someone's future behavior based on their current state of being. If you give me a sample of peoples' previous purchasing history, I can utilize machine learning to identify patterns in their shopping habits. I can then use these patterns to predict what goods someone is likely to buy next; i.e. future purchases are the unknown outcome that I want to predict. This allows me to target them with tailored promotional offers for those specific products.

Predicting the future behavior of individuals is one of the most widely used applications of machine learning but there are lots of other situations and problems that machine learning can be applied to. All you need is some unknown event or thing that you want to determine, and this could be in the past, present or future. Doctors examine their patients, carry out tests and question them about their symptoms in order to gather evidence (data). They then use this data to form a view as to what they think is wrong with the patient. They aren't making a prediction about the patient's future health but trying to work out what's wrong with them today. Doctors can do this with a high degree of accuracy because they cross reference the patient information they've obtained against what they've learnt from years of training and practice. In other words, they are looking for how

the patient's symptoms correlate with their knowledge of known illnesses. Machine learning can be applied in the same way. Given a host of information about the symptoms of different illnesses, machine learning can be used to estimate the probability that someone has a certain condition based on the symptoms that they present.

Another way to think about machine learning is as a method of reducing uncertainty. In any given situation, there are a whole host of possible outcomes that could occur. Machine learning won't tell you with absolute certainty what will happen but it can provide some insight into the likelihood, or odds, of each outcome. You may know that when someone goes grocery shopping they often buy bread, wine and chicken but, with machine learning, you can determine that there is say, a 95% chance that the next product they buy is bread, a 4% chance that they buy wine and a 1% chance that they buy chicken. Therefore, if you want them to make their next purchase in your store, you are far more likely to win their custom with a bread offer rather than a wine or chicken offer.

A **predictive model** (or just model going forward) is the output generated by the machine learning process. The model captures the relationships (patterns) that have been uncovered by the analytics process. Once a model has been created, it can be used to generate new predictions. Organizations use the model's predictions to decide what to do or how to treat people. So, machine learning is a process and a model is the end product generated by that process.

There are lots of different types of model, and there are dozens, if not hundreds, of techniques and algorithms that can be used to generate a model. However, regardless of the type of model or the mathematics used to create it, a model's predictions are almost always represented by a number - a **score**. The higher the score the more likely someone, or something, is to behave in the way the model predicts. The lower the score, the less likely they are to behave in that way.

Machine learning can be applied in all sorts of situations and to

many types of problem. However, the most common business applications of machine learning, and the ones that are the main focus of this book, relate to what people are going to do or how they will behave in the future, based on what you know about them today[3].

One, very well known, application of machine learning is credit scoring. When someone wants a loan, credit card or mortgage the lender asks the individual questions about themselves and their lifestyle. They then combine this with information from a credit report containing details about the individual's previous borrowing history, provided by a credit reference agency such as Experian or Equifax[4]. The information is then fed into a predictive model to generate a credit score.

If you live in the USA, you'll probably be familiar with FICO and/or Vantage scores. A high score (>750) is a prediction that someone is very likely to repay any money they borrow; i.e. that they are creditworthy. A low score (<500) indicates that someone is relatively uncreditworthy. Banks and finance companies the world over use similar credit scoring methods.

Another common application of machine learning is target marketing. Given information about someone's age, gender, web-browsing, purchase history, location and so on, a marketing department can predict if the person is interested in a particular product or not. They then use that prediction to decide whether or not to target them with promotional offers. Likewise, predictive models can also be used to infer how much people are willing to pay for products like insurance. This information is then used to tailor a personalized pricing strategy to each person's individual circumstances.

A further example of machine learning in action is preventative health care. Traditional healthcare systems are reactive. People seek medical assistance when they feel ill. Doctors then do their best to treat the illnesses they are presented with – treatments that can be very costly and time consuming. These days, advanced healthcare

systems are increasingly focusing their attention on prevention rather than cure. This vastly reduces costs and improves patient outcomes. Machine learning is used to assess people's medical records and predict the likelihood of them developing specific conditions such as heart disease or diabetes, often years in advance. Individuals who come at the top of the pile; i.e. those that the model predicts are most likely to get the disease, are contacted with a view of initiating preventative action. For example, making lifestyle changes or prescribing medication.

If you want to see some advanced machine learning in action, then Microsoft's CaptionBot is a good place to start[5]. Upload a picture to the website and the software will describe what it sees. I found it quite fun because it describes many images surprisingly well, but sometimes, it provided non-sensical answers or gave answers that were in the right ball park but clearly wrong. An image of someone smoking was described as: "A man brushing his teeth" and it thought that a clarinet was a guitar. It was clever enough to understand that it was looking at a musical instrument but I was at a loss to understand how it could mistake the two.

Another one to try is the thispersondoesnotexist[6] website. This generates pictures of people that look real but are completely artificial. The people in the photos don't exist.

Machine learning applications are also moving beyond just assessing information and decision-making and more into the physical world; i.e. combining machine learning with advances in engineering and robotics. This facilitates the automation of many physical tasks that would once have been undertaken or controlled by a person. Large warehouses are a prime example of a working environment that has seen a proliferation of robots in recent years with far less requirement for human staff to do things like stocking shelves or picking items for use/dispatch. If you want to see some examples of artificial intelligence combined with advanced robotics, then Boston Dynamics has lots of videos of their intelligent robots in action on their website[7].

A final example of machine learning in action is determining what type of news (and other) articles to recommend to people. Social media providers use machine learning to analyze what articles you've read in the past and the type of topics you discuss with friends. This then drives the content that they promote to you.

That's just a few ways in which machine learning is being used. Today, machine learning supports a huge range of applications. In fact, almost any aspect of life that involves decision-making in one form or another. The algorithms that match people on dating sites, facial recognition, the technology used to detect credit card fraud and systems for identifying terrorist suspects all utilize predictive models derived using machine learning.

That brings us on to the question as to what one means by **Artificial Intelligence** or AI. There are many and varied definitions of what AI is and there isn't one standard definition that everyone agrees with. However, a simple (and simplistic) definition is that a machine, that can act and reason just like a person, embodies everything it means to be artificially intelligent. This isn't an unreasonable place to start and leads one to envisage a world filled with human-like robots doing human-like things. In fact, the ability to replicate human behavior is a key feature of the famous **Turing Test** devised by the mathematician Alan Turing.

In the test, a human judge converses with a human and a machine without knowing which is which. The judge can engage them in conversation about the weather, politics, the latest YouTube clip they'd seen or anything else they liked. The machine is considered intelligent if the judge is unable to identify which is the human and which is the machine.

The Turing Test is interesting but has its limitations. In trying to pass the test, the focus has mainly been on fooling the judge rather than creating real intelligence. Those trying to pass the test haven't built a really excellent robot and then thought "I wonder if it'd be any good at the Turing Test? Hey robot, do you want to give it a go?" Instead, they've considered what they need to do to pass the

test and tried to build something that does that.

Another argument against the Turing Test is that a machine doesn't need to be self-aware like people are to pass the test. It doesn't need to understand the responses it makes. It can pass by just blindly following some very clever computer code. Some experts also question if true artificial intelligence can ever be achieved by just following a brute force approach of developing ever more complex algorithms using ever more powerful computer hardware. Is there some additional (as yet unknown) element required for human-like intelligence and self-awareness that can't be replicated via computation alone[8]?

Regardless of how one defines AI, I doubt any experts would disagree that it involves reasoning and decision-making, and for the purposes of this book, a simple working definition of AI that we shall adhere to is:

> **Artificial Intelligence (AI) is the ability of a machine to assess a situation and then make an *informed* decision in pursuit of some aim or objective.**

This definition encompasses both conscious and unconscious decision-making. An example of conscious human decision-making is where someone is deciding who would be the best candidate for a job. A machine capable of doing this instead of a hiring manager would qualify as intelligent within this area of expertise. An example of unconscious decision-making is looking at a picture and knowing that you are looking at a cat rather than a cake. Again, a machine that can do this would fall within our definition of artificial intelligence. Both of these examples of AI are in widespread use today.

In displaying intelligent behavior, a machine can make a decision on the basis of what it knows – an informed decision. If circumstances change; i.e. new information becomes available, then

the machine will reassess the situation and a different decision may result. Introduce pictures of potatoes, and the object recognition system can learn to identify not just cats and cakes but potatoes as well.

If we look at the landscape today, there are lots of clever systems that satisfy this definition of AI. When you hear about some new application of artificial intelligence, this is pretty much the definition that is being referred to.

This definition of artificial intelligence isn't one that all experts agree with, but from a practical perspective, it covers pretty much every tool, app and gadget that the tech companies describe as incorporating AI. A good AI-based application is one that can perform as well or better than the average person when faced with everyday tasks. For example, the ability to identify people from their Facebook photos, being able to assess someone's creditworthiness more accurately than an experienced underwriter, the ability to beat the best Go and chess players or being able to spot the signs of cancer on a medical scan better than an expert radiologist can.

In terms of the overall scope of AI research, machine learning is a key field of study, but there are many others. True artificial intelligence is about much much more than just pattern recognition and prediction. However, in practice, almost every AI system in real-world use today relies heavily on machine learning. Consequently, it's not unreasonable for people to be using these terms interchangeably, even if it's not technically correct to do so.

At one level, some AI-based applications can seem almost magical to the layperson. However, like most things, once you get under the bonnet the mystique evaporates. A key mistake to avoid is thinking that current AI applications are in any way intelligent in a human conscious way. Sure, they are very complex, exceedingly clever and can be creepily lifelike at times but it's all just math at the end of the day. Most experts agree, we are years away from being able to create a machine with a human-like sense of self, or which could pass itself off as human day in, day out. That's not to say there

aren't some very good chatbots out there!

All of the AI applications in use today are what the industry refers to as **_Narrow AI_**. They are very good at behaving intelligently when applied to one well defined area of expertise. However, these systems are still a long way from **_General AI._** General AI is a system that can learn and act intelligently across a wide range of environments and problems in a similar way to a person. An AI application that is used to detect tax avoidance for example, is useless at detecting the signs of cancer from medical scans. However, a person could learn to do both these tasks if they were given suitable training. In a similar vein, a system such as Google Translate is great at understanding the spoken word but wouldn't be much use when it comes to assessing if someone on a dating site might be compatible with you.

The core components that drive most AI-based applications are:

- **Data input.** This can be sensory inputs from cameras (eyes), microphones (ears) or other sources. It also includes pre-processed data such as the information captured when someone fills in a form online, details of what someone has bought using their credit card or an individual's credit history provided by a credit reference agency.

- **Data (pre)processing.** The raw data input needs to be processed into a standard, computer friendly, format before it is ready to be used.

- **Predictive models.** These are generated by the machine learning process using past experiences; i.e. large amounts of historic data. Pre-processed data for new cases is fed into the models in order to generate fresh predictions going forward.

- **Decision rules (rule sets)**. A prediction on its own is useless. You have to decide how to use it. Decision rules are used in conjunction with data inputs and the scores from predictive models to decide what to do. Sometimes these rules are derived automatically by the machine learning algorithm, but often they will include additional rules defined by human experts/business users.

- **Response/output.** Action needs to be taken based upon the decision(s) that have been made. If the decision is that someone is creditworthy, then a credit card needs to be issued. If the decision is that someone should be hired, then they need to be sent an offer letter, given a contract to sign and so on.

It's the combination of these individual components that give us the "AI."

What makes some AI applications appear so clever is the sheer complexity of the algorithms that underpin them, combined with a slick user interface to gather data and deliver the required responses in a human friendly way. Combine these components with the latest generation of industrial machinery, or integrate them into cars and other vehicles, and one has robots that can interact with their environment and engage with us in a very human-like way.

Let's begin by considering a marketing AI application for a drinks company. The app takes information gathered about individuals from social networks and feeds it into a predictive model to determine how likely they are to buy a particular brand of whisky. The system then applies a number of rules to decide if an individual should be marketed to. The rules that might exist are:

1. If the predictive model estimates that the chance of them buying whisky is more than 90% then do nothing. They will probably buy the whisky anyway.

2. If the model estimates that the chance of them buying whisky is between 1% and 90% then send them a $5 discount coupon to try and make the whisky a more attractive offering; i.e. influence the customer's behavior to increase the chance of them buying.

3. If the model estimates the chance of them buying whisky to be less than 1% then don't do anything. They probably won't buy the whisky whatever you offer them. Therefore, it's not worth the effort trying to persuade them.

So, these rules would be derived based on some type of cost-benefit analysis, where the 1% and 90% cut-offs are deemed to be the optimal level at which to trigger marketing activity. However, other business rules would also come into play, such as:

1. NEVER make an offer to sell whisky to children or pregnant women, no matter what their propensity to buy it.

2. DO NOT send offers to people with a history of alcohol dependency.

Both groups referred to by these rules will contain lots of people who would like to buy whisky; i.e. >1% chance, but from an ethical perspective targeting children or people with alcohol problems is difficult to argue for. From a purely profit-orientated perspective, marketing to children is likely to be illegal and targeting such individuals could result in a significant amount of negative publicity. These two rules are a great example of why human expertise is required to support automated machine learning-based systems, especially where systems are being used to make risky or controversial decisions about people.

One highly publicized example of the algorithms running wild is

the case of YouTube and its ad placement policy. Many large organizations withdrew their advertising from YouTube because they were found to be placing some adverts alongside material from terrorists and other unsavory sources. It was YouTube's machine learning algorithms which decided which ads to place where that had caused the problem. Consequently, YouTube had to undertake a major review of its ad placement process[9]. As a result, many months later, they decided to revert to a manual vetting process. Every video clip had to be reviewed and approved by a real person before it was included in YouTube's service which paired advertisers with popular content[10].

A reasonable question to ask is why it took almost a year for YouTube to figure out a solution? One can't be certain as to the reason, but a very plausible answer is that they spent a lot of time trying to solve the problem using a solely automated (machine learning/AI) approach, before realizing that they needed to maintain a human element in the assessment process.

The full version of this book is available in print and electronic formats at Amazon, iBooks and all good book stores.

Notes

[1] O'Neill, B. (2019). "Failure rates for analytics, AI and big data projects =85% - Yikes!" Designing for Analytics.
 https://designingforanalytics.com/resources/failure-rates-for-analytics-bi-iot-and-big-data-projects-85-yikes/ accessed 12/03/2021.

[2] Henrion, M. (2019). Why most big data analytics projects fail.
How to succeed by engaging with your clients.
https://pubsonline.informs.org/do/10.1287/orms.2019.06.08/full/
accessed 12/03/2021.

[3] Predicting consumer behavior is a very common application of machine learning but there many are others. For example, the same techniques are used to predict stock prices, when complex machines are likely to break down and which organizations are likely to become bankrupt.

[4] The original role of credit reference agencies (also known as credit reporting agencies or credit bureaus) was as a central repository for data about debts and loan repayments. This is still at the core of what they do, but these days, credit reference agencies also hold all sorts of other personal information. Consequently, a credit report can contain a wide variety of personal data in addition to information about a person's credit history. Credit reference agencies were arguably the first "Big Data" companies, decades before the term began to be applied to the likes of Google, Amazon, Facebook, et al.

[5] https://captionbotdemo.azurewebsites.net/ accessed 04/04/2021.

[6] https://thispersondoesnotexist.com/, accessed 05/04/2021.

[7] https://www.bostondynamics.com/robots, accessed 15/02/2021

[8] See, for example, the arguments made by Roger Penrose in his books: 'The Emperor's New Mind: Concerning Computers, Minds, and the Laws

of Physics' and 'Shadows Of The Mind: A Search for the Missing Science of Consciousness.'

[9] Cellan-Jones, R. (2017). 'Can Google police YouTube?' BBC http://www.bbc.co.uk/news/technology-39338009, accessed 05/04/2021.

[10] Kelion, L. (2018). 'YouTube toughens advert payment rules'. BBC http://www.bbc.co.uk/news/technology-42716393, accessed 05/04/2021.