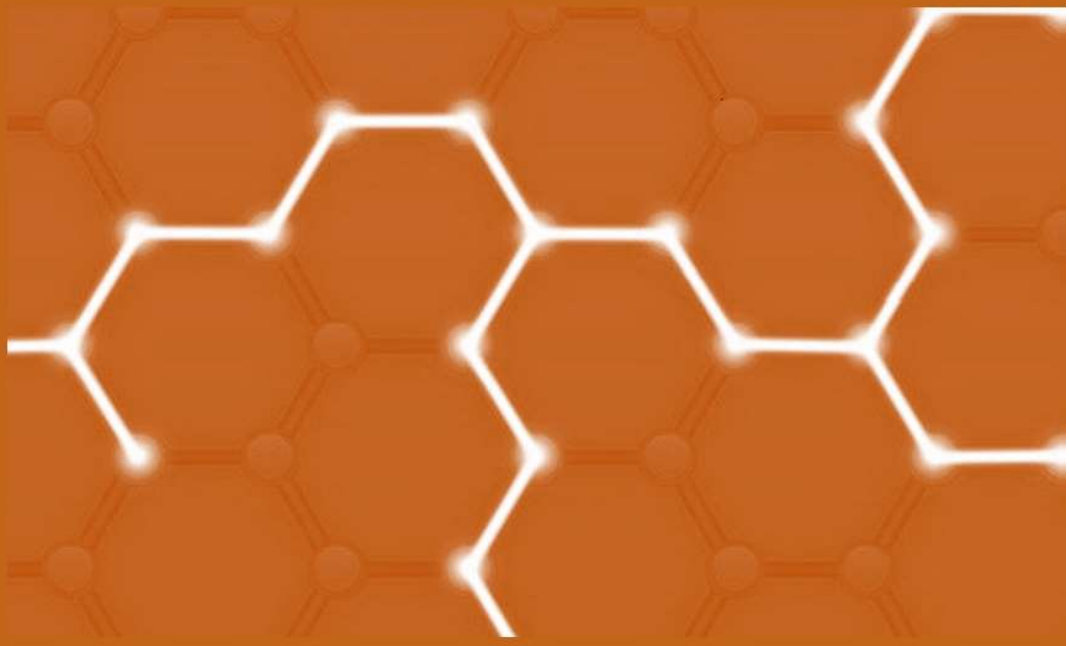

Multidisciplinary Aspects of Artificial Intelligence

Natascha van Duuren LLM
Victor de Pous LLM (eds)

Multidisciplinary aspects of artificial intelligence

Multidisciplinary aspects of artificial intelligence

Natascha van Duuren and
Victor de Pous (ed.)

ISBN: 978-90-9034111-8
NUR: 820

Contents

Foreword

While Tesla Motors CEO Elon Musk expects increasingly intelligent computers to overtake humans in the foreseeable future, with this trend potentially signalling the end of our civilisation, Alibaba Group Holding co-founder Jack Ma sees the use of artificial intelligence bringing mainly advantages. People will be able to work less, for instance, for the simple reason that systems and robots can perform the labour. Quo vadis? The Royal Dutch Association of ICT and Information Professionals (KNVI) knows in any event that from the perspective of self-learning digital applications, it is precisely the aspects of 'smart humanity' that we cannot do without, the central and simultaneously connecting theme of our professional organisation.

In order to support both our colleagues in the field and society as a whole, there needs to be more discussion on ethics. After all, there are hardly any practical impediments to information technology. It is up to the professionals who design the digital technology, work out the algorithms, apply the technology and help make information accessible for others to be more conscious than ever before about the implications of that use. What societal impact can be expected in the specific situation?

Digital professionals have traditionally played a key role in how ICT and data processing develop both in theory and in day-to-day life. This also requires a different outlook on education programmes and competency development. The broad development of competencies has become essential for the field. This is in line with the thinking of De Nederlandsche Bank as well, or so we can conclude from its guidelines (25 July 2019) for the use of artificial intelligence in financial services published as a discussion document. The regulator argues emphatically for the responsible use of AI. One of the guidelines dictates that within a financial undertaking, every individual - from the shop floor through to the board room - must have the appropriate level of skills and be familiar with the advantages and limitations of the AI systems with which they work. In our view, this is an excellent example of 'smart humanity'.

We hope this new collection of articles from the KNVI will contribute to the further development of the field among professionals. The KNVI sees the role of the Netherlands as a frontrunner in both realising AI algorithms (in advance) and in the transparent use of these algorithms (good control after the fact) as the preferred course to be taken. We also believe that the Dutch government's elaboration of the national strategy for AI with numerous parties, including the KNVI, must involve collaboration and that there should be a broad push for European collaboration. Only then will Europe, and with it the Netherlands, have any chance of succeeding in remaining a frontrunner in this area. This new collection is a contribution from the KNVI towards this ambition and provides an excellent basis for further discussion and awareness of the impact of creating and using this technology.

Wouter Bronsgeest and Paul Baak, joint presidents of the KNVI
Amsterdam, 10 January 2020

Editorial

You could say that the rise of digitalisation - the continuous development of electronic data processing - has largely taken place under its own steam, with information technology and the sector functioning as an engine and also conceiving of the road map. The role played by the government in this progress and the countless applications has been limited, despite the abundance of digital regulation in Europe over the past three decades.

This situation appears to have finally come to an end. Policy documents show that active government involvement, based on society-wide collaboration, is required to both accelerate and steer the digital developments. With forces joined, is the new axiom. Countries and governments that have been sleeping through have apparently awoken. After 75 years, time to really get to work on electronic data processing. 'Organic' growth in this domain is going too slowly, especially if a sovereign state like the Netherlands wishes to be a frontrunner.

Our government has formulated a three-part ambition: becoming a digital frontrunner in Europe, ensuring everyone participates ('inclusivity') and realising a good basis of trust for the digital economy, government and society. This is evident from the Dutch digitalisation strategy 2.0 of 5 July 2019. While ICT was suddenly labelled the 'axis of innovation' at the beginning of this century, today it is invariably referred to as 'key technology'. Not exclusively for ICT generally, but particularly in relation to artificial intelligence (AI) and other 'transformative' digital technologies, such as blockchain. According to the latest time line, the government's Strategic AI Action Plan is due to be published in autumn 2019 (after this collection has been put together). The crux of this policy document is clear: the further development of *artificial intelligence as a national priority*.

AINED, a public-private partnership in our country of research institutions and employers, has already made a move and in fact regards AI as an essential condition for *safeguarding* our prosperity and international competitive position. But the general common thread in relation to AI is, from an international perspective, less strict and more cautious in nature. Artificial intelligence has the ability to *improve* our well-being.

This collection aims to contribute to the insight into artificial intelligence by clarifying the much talked about and transformative key technology from various angles, as was done earlier in relation to blockchain.[1]

After all, AI is a perfect example of a topic that requires a broad-based multidisciplinary approach. It is important to approach the countless issues and challenges related to (the use of) AI in our society, not only from a legal or technical perspective. Education, digital skills, governance issues, etc play an important role as well and will even be determinative for the further introduction of AI in our society.

The urgency of tackling the issues relating to AI quickly and from a broad-based perspective is frequently emphasised in the media. Following on a survey among CEOs from 37 countries, it emerges that businesses are lacking both the employees and resources for the AI era.

In our country alone, the researchers say that approximately 3 million people need additional training because AI will drastically change the way we work. This is a worrying message, certainly given the fact that the Netherlands actually wants to build up a head start in terms of AI knowledge.

The Royal Association of Information Professionals (KNVI) is the ideal party to contribute to the Netherlands keeping up in the race for AI. After all, the KNVI is *the* platform for digital professionals from various segments of society that can shed light on a topic like AI from a broader perspective.

We thank all the authors who contributed to this collection. The contributions were written in a personal capacity and were deliberately kept concise in nature. The chapters contain valuable analyses and suggestions, but emphatically do not provide advice for concrete cases. They are based on current knowledge and the current state of affairs. We are thoroughly aware that developments in AI are taking place at lightning speed and that this will inevitably mean that certain insights will be outdated in the foreseeable future and will require revision.

Nevertheless, we are convinced that this book will contribute to multidisciplinary knowledge sharing on AI and that it is important that the work is being published at this time. Knowledge-sharing that is not only relevant for those already working with (concrete applications of) AI now or who will be doing so in the near future. This collection could also be valuable for policy-makers, including

[1] N.H.A. van Duuren and V.A. de Pous (ed.), *Multidisciplinaire aspecten van blockchain* [Multidisciplinary aspects of blockchain], Amsterdam, 2019.

the Dutch AI Coalition, precisely because the opportunities but also the challenges of AI are outlined from a broad-based practice.

Natascha van Duuren, Victor de Pous Amsterdam, 10 January 2020

1. The future is today

Wiemer Kuik and Wim Kweekel

Investments in artificial intelligence are growing rapidly, and the number of applications is increasing substantially. The McKinsey research agency[1] estimates that twenty to thirty billion dollars is being spent on AI worldwide. Ninety percent of that is going to research & development and implementation. AI-fed systems can already fly independently (drones), beat human game players,[2] translate texts, drive cars, trade shares, develop new pharmacological treatments, detect cracks in aeroplane engines, discover planets and much, much more. Multinationals like Facebook, Google, Amazon, Apple, Microsoft, IBM and Baidu are major investors and frontrunners in the development of AI. But the ICT sector as a whole is also focused fully on this disruptive technology, which entails major changes for our society. AI use today is still just in its infancy.

What is AI?

The application of artificial intelligence (AI) can already boast impressive results. Take healthcare, for example[3]. Technology is already being applied in the early detection of dementia, predicting lung cancer, predicting patient outcomes in intensive care, detecting abnormal cells in pathology and, for instance, identifying dermatological abnormalities. All involving extremely valuable analyses by computers that use a great deal of data and then translate these to individual patients.

Artificial intelligence is the area of informatics that focuses on systems that perform functions we normally associate with the human brain, such as learning, dealing with ambiguity, solving problems, recognising emotions and even being

[1] https://www.mckinsey.com/~/media/McKinsey/Industries/Advanced%20Electronics/Our%20Insights/How%20artificial%20intelligence%20can%20deliver%20real%20value%20to%20companies/MGI-Artificial-Intelligence-Discussion-paper.ashx

[2] See the documentary on Netflix https://www.netflix.com/nl/title/80190844.

[3] https://www.nictiz.nl/wp-content/uploads/Rapport_artificial_intelligence_in_de_zorg.pdf

creative.[4] In doing so, AI uses rules formulated by humans, or learns itself from sample data. This information technology has a long history. The first generation of successful AI systems, which reached its peak around 2000, worked on the basis of explicitly formulated rules. We call this knowledge technology. Being able to set up a cohesive set of knowledge rules is an essential condition for knowledge technology. An algorithm uses these rules to figure out the most likely solution to a particular problem, for example the shortest route to a destination or the most likely explanation for a set of symptoms.

We are now in the prime of the second generation: machine learning. Here the system does not need any preformulated rules, but instead learns itself from large quantities of sample data what the correct answer, solution or way of doing something is. AI is now able to itself deduce the rules from the data available.

Applications using AI

Artificial intelligence has long been a part of our life. Take the computer systems behind Spotify and Netflix, for instance. They monitor the user's choices and in doing so learn the user's preferences. This enables these streaming services to recommend to users music and films that match their personal taste. Computer games are intelligent as well. In first-person shooters, the virtual enemy analyses his surroundings to be able to seek cover. Online shops use artificial intelligence to predict purchasing behaviour and to show the customer relevant products.

Apple, Google, Amazon, Microsoft and others use AI in their virtual personal assistants. And self-driving cars are able to make decisions based on algorithms and to improve their own driving abilities based on driving experience. AI is also used to write news reports on companies' quarterly figures, identify objects, navigate robot vacuum cleaners, translate conversations and predict home burglaries. Companies use this technology for their online customer service with chatbots, to optimise processes, for their inventory and supply chain network, to predict current commodity prices and demand for products. These are just a few examples from our day-to-day life. Artificial intelligence is not some futuristic notion, therefore. We are already dealing with it.

[4] Definition, as contained in https://www.nwo.nl/documents/enw/rapport-ai-voor-nederland-vergroten-versnellen-en-verbinden.

Growth in the types of application

The initial applications were in the world of data.[5] It started with the structured analysis of data in systems in order to perform advanced analyses. Then applications arose that analysed our online behaviour and provided us with recommendations, enabling companies to micro-target customers.

With the Internet of Things (in which all sorts of sensors register data), new and rapid analyses could already be made, so that maintenance became predictable, for instance. These days AI can not only be found in the world of data, but also in the 'real' world. Patterns are recognised and human behaviour can be predicted. And last but not least, it can support and assist people through natural language and speech recognition. As such AI is becoming less data and computer-focused and increasingly people-focused: human-centric AI.

Opportunities and possibilities

AI offers many opportunities and possibilities for businesses and government organisations. Using algorithms to analyse data can and will result in cost savings.[6] In production and also in management processes, for instance.[7] On the other hand, AI can be used for all sorts of contact moments with customers/users. From the use of speech recognition in Amazon's echo, for example, to how we have contact with computers, systems and robots. AI is used in that case to analyse large volumes of use data and consequently realise an even better user interface.[8]

Data-oriented applications

AI is based on data. The more and faster data is available, the better algorithms can do their work. In addition to AI, there are a number of other developments that will accelerate the collection, dissemination, analysis and processing of data.

To start with, the further spread of the Internet of Things (IoT), which will mean sensors at many more locations will be recording data. An unbelievably important source that did not previously exist. Then the new 5G network, which is

[5] https://www.bcg.com/publications/collections/strategy-digital-artificial-intelligence-business.aspx

[6] https://www.bcg.com/publications/2018/artificial-intelligence-will-reshape-companies-industries-nations-interview-kai-fu-lee.aspx

[7] A fine example: https://fd.nl/achtergrond/1290746/de-nieuwe-terreur-van-de-moderne-manager-heet-hyper-awareness

[8] Another example is application in the financial sector. For example, see: https://www.solidprofessionals.nl/wp-content/uploads/2019/02/2019018-01-A4-artikel-AI-def-LR.pdf

expected to allow the sending of 50 times more data during the same time span as currently. Finally, there is the phenomenon of edge computing, which involves smart interaction between local computers (such as smartphones, for instance) and Cloud processing so that more tasks can be efficiently performed.[9]

All of this will generate many brand-new applications. Imagine what we will be able to do in education if we record the classroom activities. Or in airport security, if we can detect potential terrorists among the ordinary passengers. As a follow-up to this, we now see what is referred to as autonomous AI. Data are analysed and fed back to machines such as robots and vehicles that can then function autonomously, without human intervention. Autonomous AI will turn the world of existing products, businesses and revenue models entirely on its head. This will force businesses, sooner or later, to rethink their products and services, their revenue model and their production and aftercare processes.

Impact

According to the Dutch government, of all new digital technologies, AI is expected to have the greatest impact on the economy and society in the coming decade. AI is developing rapidly and is being applied in virtually every conceivable sector and domain. The Dutch government sees possibilities for AI in virtually every corner of society, education, healthcare, agriculture, mobility, energy, industry, small and medium-sized enterprise and the government itself.

Applying AI requires a great deal of funding and knowledge. But as has been the case for most of the technological developments over the past 20 years, 'the winner takes it all'. In other words, the first party to have and be able to use a good AI algorithm will take over many markets.[10] 'The winner takes it all' applies not only for businesses, but for countries as well. China is currently dominating the AI world.[11] 85 percent of Chinese businesses use AI in their own company. China has the ambition of being market leader by 2030 and is emphatically taking the initiative. The US does not want to be left behind and is increasingly making AI

[9] Such as our telephones, for example. An iPhone today has more computing power than the computers used to launch the Apollo space missions. An iPhone uses AI in facial recognition, incidentally.

[10] Some experts reason that the best AI application will ultimately be the last human invention. Because of the rapid self-learning ability. This is known as 'intelligent explosion'. Also see: https://becominghuman.ai/in-the-battle-for-artificial-intelligence-winner-takes-all-4c2447e68237

[11] https://www.forbes.com/sites/louiscolumbus/2018/12/16/how-china-is-dominating-artificial-intelligence/#2b7e77be2b2f

research a priority.[12]

This is a worrying trend for a country like the Netherlands. The NRC rightly devoted an editorial feature to the topic. The title is very telling: 'Nederland mag geen AI-kolonie worden van de VS en China' [The Netherlands must not become an AI colony of the US and China].[13] Everything and everyone will try to be the winner. With AI, the possibilities for improving society are unlimited. At the same time, there is clearly a major risk of unintended consequences and malicious parties; just a mention of the company Cambridge Analytics probably suffices. Because AI can also put important fundamental rights and public values under pressure, such as the prohibition against discrimination and the protection of human dignity and autonomy.

The Dutch perspective

AI is developing at lightning speed and the countries around us[14] are trying to secure a spot among the frontrunners. They are doing this with a national strategy and substantial investments. Since AI is a generically applicable technology, it will affect nearly every aspect of our economy and society. The 'winner takes all' effect makes it even more important for the country itself to get to work on AI. In the recently published report 'AI voor Nederland' [AI for the Netherlands], it is investigated to what extent the Netherlands is capable of taking serious steps in the area of AI.[15] Some conclusions.

— **Business sector.** The majority of Dutch businesses recognise the potential impact of AI on their own activities and say that the use of AI must be a priority. AI activities are still in the planning or pilot phase at most businesses. AI adoption among SMEs and the number of start-ups and scale-ups must increase substantially. On average, the business sector can be aided with more knowledge, more talent, more suitable data and clear and helpful regulation.

— **Education.** There is strong growth in the demand for places in Computer Science and AI-related degree programmes at universities and in higher

[12] https://tweakers.net/nieuws/148972/vs-geeft-onderzoek-naar-kunstmatige-intelligentie-prioriteit.html

[13] https://www.nrc.nl/nieuws/2018/08/28/artificiele-intelligentie-nederland-mag-geen-ai-kolonie-worden-van-de-vs-en-china-a1614472

[14] https://www.rvo.nl/sites/default/files/2018/10/Artificial-%20Intelligence-%20Holland-Innovation-Network-Special.pdf

[15] https://www.nwo.nl/documents/enw/rapport-ai-voor-nederland-vergroten-versnellen-en-verbinden

professional education. The current capacity is still not adequately attuned to this. We see that fewer and fewer secondary schools offer Informatics as a subject because of a shortage of suitable instructors. The Netherlands is currently lacking a system-wide approach for the increased demand for AI-related programmes at all levels of education.

— **Research**. The quality of Dutch Computer Science research is excellent and among the best in the world; AI already played a leading role in the 1970s. There are also public-private partnerships in AI in various ecosystems. At the same time, scientific production is on the decline (relatively), from approximately 3 percent of world output in 2011 to less than 1 percent in 2017. With a clear choice for more public-private research into AI and the availability of adequate research and other funding, the Netherlands can significantly strengthen its position in an international context and its appeal to AI talent.

— **Social domain.** While some jobs and perhaps businesses will disappear over time, other jobs and businesses will be created. This will require educational institutions, employment agencies and workers themselves to adapt. There is increasing societal and political attention to the ethical aspects of AI because equal treatment, autonomy and human dignity could come under pressure. Given the speed at which AI is developing, we will have to dynamically shape policy and regulation in parallel to the development of AI.

Conclusion

Artificial intelligence is not a revolution that is suddenly happening to us; it is a development gradually making its way into our society, often without us noticing. Little by little, AI is increasingly becoming a part of our day-to-day existence. AI brings with it a great many possibilities. The businesses and states that are the first to take major steps in developing and applying AI have the opportunity to avail of the 'winner takes it all' principle, with all the consequences this entails for other businesses, industries and even countries.

To continue to keep pace globally, investments are needed in knowledge, infrastructure and research. A role is reserved in this for both the government and for the business sector. Businesses, government, knowledge institutes, politicians, administrators and end users will have to work together in a timely fashion to ensure that important societal and ethical values are safeguarded and taken into account in the development of artificial intelligence. No one benefits from a society that makes large-scale use of intelligent systems without any sense of ethics.

Points for attention

- AI poses both an opportunity and a threat. There are possibilities for using AI in every branch of industry. Those who take the first steps will have the best possibilities. Those who fail to make a start run a serious risk when it comes to their continuity and right of existence in the long term.

- Because of the major societal impact of AI, it is important that businesses, government, knowledge institutes, politicians, administrators and end users engage in dialogue on this topic. How will we make the most of the opportunities afforded by artificial intelligence?

- AI involves a number of particular ethical and legal issues. Discussions and measures are inevitable, at least if we do not want to leave it to a number of large parties (such as Facebook, Google, etc) and countries, including China and the US.

2 From programmed logic to machine learning

Sieuwert van Otterloo

Programmed logic is computer technology that we have been using for decades: rules laid down in code to process and evaluate data. Well-known examples of this include spell checks that underline unknown and misspelled words or websites that calculate shipping costs for an order. Machine learning, on the other hand, is a collective name for all sorts of new algorithms coming from AI research, which are used to solve problems that cannot be solved with programmed logic, or only with great difficulty. Take, for instance, reading handwriting, recognising people in photographs, giving advice on what products to buy or films to watch, providing advice via chat or vehicles that drive autonomously. These are all cases in which machine learning algorithms work much better than programmed logic. Businesses are therefore rushing to use these to improve their service provision and cut costs. Google, Microsoft and others are anticipating this by offering machine learning as a service via their Cloud platforms.

Why is machine learning different?

Programmed logic is not always correct, but it is always consistent and reliable. People are now accustomed to using logic and have adapted to it. Organisations have done this as well. They have set up test processes for IT systems and have support and contact options where people can report errors. These errors are then reproduced and the logic is adjusted until the outcomes are correct.

What many businesses do not fully realise is that machine learning is not automatically better, but is above all very different. Machine learning algorithms must be created and tested in a different way. If you fail to do this, you run many new risks.

The reason why machine learning algorithms work differently is because the result depends entirely on the data used to train the algorithm. After the algorithm has been developed and tested by the supplier, the algorithm needs to be trained using data from practice. The system (algorithm plus training data) then needs to be tested again using test data. This gives a picture of the strengths and weaknesses

of the algorithm, after which one must decide whether and how to use the algorithm.

Many organisations underestimate these extra steps and perform them carelessly. This results in algorithms being used in situations in which they work poorly, without the users and organisations realising this.

Possibilities and advantages of machine learning

In the right circumstances, machine learning can be faster and more precise than programmed logic or human experts. This is because machine learning can be based on more data and can quickly adapt to new data. Examples from practice include:

- Recognising rare diseases. During their careers, individual doctors will encounter certain illnesses more often than average and others not at all. This can result in bias, whereby certain illnesses are diagnosed too hastily and others are not recognised. An algorithm can be trained using the experience of many doctors and can therefore better estimate the likelihood of less common symptoms.
- Integration of video and audio. With good implementation, the outcomes of machine learning are integrated, whereby machine learning looks at the total media data and can also generate images and sound. This is valuable for automatic interpretation or translation, for instance; the translated text is inserted immediately. To see how quickly that works, take a look at this music video from Google ('La Bamba'), in which text is recognised and automatically translated.[1]
- Learning from feedback. Video services like YouTube and Netflix provide people with personal recommendations. These recommendations are adaptive: the platform keeps track of what viewing tips you actually follow and adapts the recommendations accordingly. Many people feel that the recommendations get increasingly better as a result. This even works if your feedback is not explicit; by just measuring whether you watch a video to the end, the algorithm already knows whether certain videos appeal to you.
- Discovering weak connections. Machine learning is much better at discovering weak connections than humans are and can therefore use a large quantity of variables that are combined in larger models. It is therefore possible to predict

[1] https://www.youtube.com/watch?v=06olHmcJjS0

borrowing capacity, for example, based on someone's overall payment behaviour and income history instead of just their current income.

Because machine learning can sometimes work unexpectedly well, there are also businesses that share their data externally and invite other machine learning experts to present alternative models. The platform 'Kaggle' (https://www.kaggle.com) contains datasets and questions from many businesses. There is often even a reward for a better solution, for instance for predicting stock market prices based on news reports, discovering oil from geographic data, improving the texts of job vacancies and a great deal of other data as well. In all these cases, organisations believe that machine learning can discover connections that the organisation has not yet figured out itself using traditional logic.

Risks of machine learning in practice

The fact that even large businesses underestimate the development of machine learning that functions well is evident from the reputational damage Microsoft suffered as a result of the development of the Tay chatbot.[2] This chatbot was developed and tested by Microsoft and then made available on Twitter, where anyone could chat with this chatbot. The chatbot used all the conversations conducted as new training data, however, which meant that the chatbot changed as a result of use. Microsoft probably expected that this machine learning would cause the chatbot to get better and better. In practice, users fed the chatbot racist texts. The chatbot reproduced these racist statements and therefore 'became' racist itself.

In this example, you see clearly that the chatbot's behaviour changed as a result of the influence of data. The makers clearly underestimated the resourcefulness and wantonness of internet users.

There are also many other examples of machine learning applications that did not work well. For example fraud detection algorithms for which it cannot be verified whether they work or not, or undesirable and offensive recommendations. The interplay of data and algorithm causes unexpected behaviour in each of these cases. The idea that an application is reliable after it has been properly tested applies for programmed logic, but not for machine learning.

[2] https://www.bbc.com/news/technology-35902104

Dealing with data

When using and evaluating machine learning, a great deal of attention must be paid to the data used. That starts with transparency. It must be clear which dataset was used to train the algorithm and which dataset was used to test. The question is then whether these training data and test data are representative for use; are the examples that the algorithm receives comparable to the data it will receive later? Algorithms for a self-driving car that has only been trained and tested for dry weather cannot be relied on for use in rain or snow, for instance. It must also be known whether the algorithm is constantly trained on the basis of the latest data; in that case, constant testing and evaluation will also have to take place to check whether the system is satisfactory.

Due consideration must also be given to the ownership of the data used. When a machine learning cloud service is used, the data is passed on to the cloud provider. Is the provider allowed to use the data to improve performance for other users? Can the data be reproduced by the algorithm, allowing sensitive information to be leaked? And do the data include personal data? In that case, the new GDPR privacy legislation applies and extra precautions must be taken.

Evaluation of performance

With programmed logic, it is customary to test until the logic is correct for all known test cases. This creates a large degree of certainty about the correct outcomes of the system. You might expect the same with machine learning, but unfortunately the situation is different in practice. The performance on the test dataset is often not perfect, but rather 'good', 'reasonable' or even 'moderate'. It is then up to the machine learning specialist to take a decision whether to continue training or start using the system. The decision is often taken to just start using the system, despite the test results. This is because it is often difficult to get the extra data needed to further improve the system, and because the machine learning system is considered 'good enough'. This decision that it is 'good enough' is often taken too hastily, unfortunately. In many cases, only the average performance is looked at and not the outliers, for instance. A good example is facial recognition.[3] Known algorithms do this reasonably well (90% correct recognition) on the entire dataset. If you look at minority groups, however (women, people with darker skin), the performance is much worse. It seems that there is not enough diversity in the training data used.

[3] http://gendershades.org

It is important that the tests performed and the scores achieved are disclosed so that everyone becomes more aware of the fact that machine learning outcomes are not always correct. The organisation that uses the outcomes will then have to look into whether adequate tests have been performed using realistic and diverse data. If the score is not perfect, this could have implications for the application: you will have to inform users that the advice could be incorrect and you may have to include extra manual checks in the process.

Conclusion

When using machine learning, there is a greater likelihood of practical problems than when using programmed logic. Extra measures are therefore also needed for responsible use. A great deal of attention will have to be devoted to using the correct data for training an algorithm: there is no machine learning algorithm in the world that will perform well if it has not been trained using the correct data. What the correct data are differs from application to application. After training, extra testing of the algorithm will also have to take place using appropriate test data.

It is then possible that the machine learning system will develop in practice. With machine learning, it is therefore necessary to continue to evaluate the performance. Since both machine learning and programmed logic can be used, it is very important that users know which of the two they are dealing with. This is the only way for users to know that the outcomes might not be correct and that they can confront the organisation about incorrect outcomes and ask for a reassessment, for instance.

Points for attention

- With machine learning, always find out what data were used for training and testing.
- Be transparent about the use of machine learning.
- Take into account that machine learning applications involve a margin of error and make sure there are appropriate checks in practice.

3 An overlooked handicap of the head start

Ken van Ierlant

Artificial intelligence has become the new buzz word and a hype in the digital world. University courses on this subject have been fully booked for years, but what does it entail exactly? AI was first defined in 1959 by the American professor John McCarthy.[1] '*Every aspect of learning or any other feature of intelligence can in principle so precisely be described that a machine can be made to simulate it. An attempt will be made to find how to make machines use language, form abstractions and concepts, solve kinds of problems now reserved for humans, and improve themselves*.' In his vision, machine learning functions as a solution for usually complex, heterogeneous issues for civil and non-civil applications. These days, however, many organisations are lagging behind technologically, partly because of the use of legacy systems. This inhibits developments in this domain. The trick is to step up the phasing out of the application-centric systems and transform these into a data-centric 'next Gen platform'[2] [3] as a primary condition for starting to work with artificial intelligence.

AI revisited

Sixty years ago, scientists defined seven key areas for the applicability of AI. The initial areas are:

1. *Simulating higher functions of the human brain;*
2. *Programming a computer using a general language;*
3. *Arranging hypothetical neurons so that they can form concepts;*

[1] John McCarthy: Computer scientist known as the father of AI.
https://en.wikipedia.org/wiki/John_McCarthy_(computer_scientist)

[2] *Platform Theory, Technology Platforms:* https://sloanreview.mit.edu/tag/platform-theory/

[3] https://www.google.nl/url?sa=t&rct=j&q=&esrc=s&source=web&cd=7&cad=rja&uact=8&ved=2ahUKE wiu-I3MvKXkAhWF_qQKHSiqDoUQFjAGegQIChAB&url=https%3A%2F%2Fwww.crunchbase.com%2Forganization%2Fnextgen-platform&usg=AOvVaw0cIW42STlPFGNrNYBY9wJH

4. *A way to determine and measure problem complexity;*
5. *Self-improvement;*
6. *Abstraction: Defined as the quality of dealing with ideas rather than events;*
7. *Randomness and creativity.*

In the meantime we are also researching new areas where AI can be applied, such as pharmaceuticals, with DNA-related medicines, or self-driving cars and nanotechnology, for instance, but also in applications that influence consumer behaviour, resolve legal issues or predict share prices.

Realistically, we have made some progress over the past 60 years on just three of the seven aforementioned areas, specifically: computer language, self-improvement and solving complexity. Because of the rapid development of microprocessors and the much better software, the development in AI has sped up. With the possible rapid rise of quantum computing,[4] an exponential acceleration will take place in the applicability of AI and the throughput and storage of gigantic quantities of data. Experiments are now ongoing with AI in vertical data-centric environments that are usually processing-intensive. In this context, relatively small quantities of data are serving as the fuel for experiments in the medical field, but also for facial recognition, autonomous driving, security, robotisation, etc.

The Art of Computing

In computing, we have for years been distinguishing between compute and I/O intensive[4] applications.[5] Processing-intensive (compute) applications are often found in science and at NASA, for instance. In order to accurately predict the landing of a Mars rover down to the millimetre, relatively small quantities of data are recalculated iteratively for a long time. The CPU spends relatively more time on the calculation than on the transport of large quantities of data such as in business applications like ERP or CRM.

I/O-intensive applications are more focused on short CPU calculations with gigantic quantities of data that are transported from A to B.[6] The internet has become a massive data bus in which millions of computers 'pump' petabytes per

[4] https://en.wikipedia.org/wiki/Quantum_computing

[5] CPU Bound vs. I/O Bound | Computer Basics.

[6] https://www.google.nl/url?sa=t&rct=j&q=&esrc=s&source=web&cd=4&cad=rja&uact=8&ved=2ahUKEwiB5ISYvKXkAhVDhqQKHZtGBc4QFjADegQIAhAB&url=https%3A%2F%2Fwhatis.techtarget.com%2Fdefinition%2Fcompute-intensive&usg=AOvVaw0yB2oA49cfLCAOTBJfRSW2

millisecond around the world. The quantities of data are also increasing exponentially in the meantime. This boisterous growth in data is increasingly at odds with the possibilities of AI and the existing infrastructure or lack of a clear architecture.

Data-centric versus application-centric[7]

Many scientific applications that use data-centric IT landscapes are built around a vertical focal area, such as the exploration of oil and gas fields or the data from self-driving cars. When it comes to commercial applications, 99% of these have been developed in an application-centric manner. This makes predicting consumer behaviour, for instance, much more complicated because these applications - in which the data are fragmented across multiple databases - are usually linked to outdated business logic. It is as if you have thousands of aquariums, with a goldfish in every fish bowl, and 30% have mutated or already died but in the meantime continue to expand unchecked.

In other words: the big problem with AI in commercial applications is that the application-centric IT landscape that roughly all businesses - but also governments - have is not at all suited to AI. We are not yet familiar enough with the interaction between I/O-intensive and processing-intensive processes and the extent to which AI can be applied.

Two steps forward, one step back

In the past, IBM invested in the Deep Blue chess computer to demonstrate that a computer could beat a human at chess.[8] This form of 'weak AI', in which thousands of computer calculations were made before a piece could be moved on the board, was extremely processing-intensive. It yielded amusing results but cannot be used in practice for commercial applications. It simply takes too long!

These days IBM is experimenting with supercomputer Watson as a much stronger form of AI. Human Reasoning is emulated in rule-based environments like accountancy or the administration of justice, making use of extremely fast processors and modern storage and data-processing technology.[9] The outcomes of this are very promising! Presenting legal questions whereby legislation, case law

[7] The Data-Centric Revolution: Data-Centric vs. Data-Driven.

[8] https://en.wikipedia.org/wiki/Deep_Blue_(chess_computer)

[9] https://en.wikipedia.org/wiki/Watson_(computer)

and interpretations thereof are taken into account can, for instance, result in a much faster and more efficient handling of a problem.

Because most law and accountancy firms have large application-centric IT landscapes that have been cobbled together over the past 30 years, there are a lot of data that have not been categorised and that have been stored fragmented in various different data models/databases. In other words: there is a great deal of unstructured data hidden in legacy applications. Letting AI applications loose on this 'old junk' generates a disproportionate quantity of - unnecessary - data traffic, causing computers and networks to become backed up and then deadlocked. Prolonging this situation also means that the performance of existing systems deteriorates drastically and the costs of maintenance rise exponentially. It should be added to this that the outcomes of the calculations are increasingly unreliable because over the years, approximately 30% of all the data have become corrupted. We already see that in the calculations and predictions for the climate. Scientists are increasingly doubting the outcomes of their models, with the major political and social unrest that ensues.

Even aside from the technical problems, law and accountancy firms have a business model focused on maximising hourly billing. The impact of AI - if used properly - will sweep this entire sector from the market in a few years. The resistance this will provoke in these sectors because computers will be taking over the work of lawyers, judges and accountants will significantly curb the development and application of AI in this market.

Accountability

With the introduction of the General Data Protection Regulation (GDPR) as the new European privacy law, strict requirements are imposed on the processing of personal data. Failing to satisfy the 'notion of accountability' stemming from this new legislation confronts businesses and governments with major dilemmas. Because the architectures of IT landscapes have grown organically and were never designed to be put online, a situation arises in which AI-like applications, set up with the best intentions, subsequently leave behind a trail of destruction in these legacy-based environments.

The privacy legislation is a reaction to a consumer-centric society in which the individual's data must be protected. With the arrival of AI systems to predict consumer behaviour, glorifying corporate turnover and profits, this gives rise to a

very particular paradox. Profiling[10] as a new silver bullet to increase conversion is prohibited by GDPR, unless..., but the tension that these calculations and transport of data cause in the IT landscape will most likely result in data leaks and many unknown IT incidents.

The elaboration of new technology in combination with old application landscapes that have evolved incrementally is still unknown territory. The legislation may indeed be ahead of things here but without realising what requirements it is imposing for innovation and economic feasibility.

Catch 22

The moral of this story is that AI creates high expectations, and with good reason. But it could be that these are too high, because we are not at all ready for them. Despite the fact that we have seen amazing developments over the past sixty years which will now result in new forms of mobility, communication and business models, at 99% of businesses and governments the business operating models are seriously outdated and no longer satisfy the condition that new technologies be able to be applied.

What is missing is an unequivocal data-centric architecture that can handle both compute and I/O intensive applications at the same time, in an efficient, technically error-free, legally sound and above all economically feasible manner. In order to build the bridge between the applicability of AI in the scientific laboratories and the day-to-day use by consumers or individuals, a great deal still has to happen to make it economically profitable.

Conclusion

Artificial intelligence is becoming one of the many digital drivers behind new business models that will be disruptively turning sectors upside down. In numerous sectors, experiments are already ongoing with very promising results. Nonetheless, we must be cautious about raising expectations too high because, on the one hand, the power of AI rapidly jeopardises the interests of institutions and, on the other, practically all organisations are lacking the technological basis. After a run-up of 50 years, the technology and applications have suddenly leapt ahead of the market

[10] https://ec.europa.eu/newsroom/just/document.cfm?doc_id=48850 and https://www.google.nl/url?sa=t&rct=j&q=&esrc=s&source=web&cd=4&cad=rja&uact=8&ved=2ahUKEwj428XevqXkAhWLM-wKHXbdD90QFjADegQIBhAB&url=https%3A%2F%2Fwww.merriam-webster.com%2Fdictionary%2Fprofiling&usg=AOvVaw1LbMwCcp6wXMT1KaruNwmE

exponentially. This catching-up process and change in the mindset at organisations will certainly take time.

Points for attention

- Artificial intelligence is an interesting area of development with high expectations. Via this information technology, the predictability of failures can be prevented, for instance, which can subsequently prevent incidents and save lives and a great deal of money

- Expectations are tempered by the technological arrears in society.

- With AI, we see that numerous business and operating models are being disruptively undermined, with the result that carefully planned business strategies become outdated in no time.

- Technological acceleration results in AI developing faster and time will tell how quickly the achievements will win over the passive resistance.

4. New skills

Liesbeth Ruoff-van Welzen and Bramjan Mulder

Digitalisation is making its way into our working lives at all levels. Founder Klaus Schwab[1] of the World Economic Forum wrote in The Future of Jobs (2018) that the contours of a new world of work are quickly becoming visible in the fourth industrial revolution and this will require new skills from everyone. Organisations need 'digital leadership' to optimally utilise new technological possibilities. For all working people, digital skills are important to be able to keep pace with the labour process. IT professionals are expected to be able to develop systems that build a bridge between existing, old and new technological possibilities. This applies especially for AI. Because it assumes that our intelligence will be replaced or supplemented by an artificial or digital form.

Current situation

New technologies such as big data, artificial intelligence and cloud computing offer possibilities for different ways of organising and performing processes within an organisation and approaching markets. This requires a mind open to this and with the knowledge and experience to apply new digital possibilities. How do you get these developments to work for you? How do you get AI to work for you? How do you apply these developments within the limits of your own organisation? Do external employees have the right skills? Can leaders and managers monitor that the decisions they take are supported correctly by their AI systems? After all, this assumes that they can assess the system and the impact of this system on their organisation.

Jim Stolze[2] saw in 2018 that the Netherlands was falling behind in relation to AI. His column 'In Helsinki weten ze meer van KI' [They know more about AI in Helsinki] in the FD caused a buzz and prompted efforts to get the first National AI course[3]

[1] https://www.weforum.org/reports/the-future-of-jobs-report-2018

[2] https://www.jimstolze.nl/weblog/2019/01/03/ai-cursus-de-bedankjes/

[3] https://www.ai-cursus.nl/

on track to go live within a year (January 2019), based on the Helsinki model. People of all ages can get a sense of AI.

The National Register and Nederland ICT recently carried out a 'baseline measurement'[4] focused on the question of how management, boards of directors and supervisory boards view or deal with current issues relating to digitalisation.[5] That survey indicates that a majority consider ICT extremely important for their core business and that they are entirely dependent on ICT, certainly on the operational level.[6] But also that finding and deploying the right people who know their stuff are the key obstacles standing in the way of carrying out a digital transformation.[7] From the knowledge questions in this survey, it emerged that the lack of in-house knowledge at the organisation is regarded as a shortcoming (by approximately 90%) and that they are consequently dependent on external expertise.[8]

Human resources

A professional group that should certainly have an eye for the changes in the working landscape caused by digitalisation is HR professionals. A survey[9] shows that just twenty percent of this professional group has a clear picture of the effects of these technological developments on matters such as staffing, competences, recruitment, training requirements and work processes. The majority (47 percent) of the respondents acknowledge that technological developments are currently prompting major changes in the production process, products or services. Another 43 percent think that technological developments, including AI, are only causing small changes in the workplace.

Lack of understanding?

Citizens, on the other hand, are quickly becoming accustomed to the advantages of AI. The smart assistants of Google, Apple and Amazon are developing rapidly.

[4] Baseline measurement on digital transformation in board rooms in the Netherlands by Prof. Marcel Thaens and Prof. Valerie Frissen, carried out in 2018 and published in 2019.

[5] Page 10 of the Baseline measurement

[6] Pages 32 and 33 of the Baseline measurement

[7] Page 45 of the Baseline measurement

[8] Pages 52 and 53.

[9] Berenschot research - https://www.hrpraktijk.nl/topics/arbeidsvoorwaarden/nieuws/privacy-en-robotiseringdigitalisering-opvallend-lage-prio-bij-hr

They get better and better at understanding what we say, give better answers and are furthermore integrated in other programs. Photos can be searched for people, chairs and cats and our personal Facebook bubble is becoming scarily more precise. Thanks to the intuitive user interface, we no longer need to read through any manuals: how the application works becomes clear playfully. But how much understanding does the average citizen have of these digital developments?[10] Is he or she aware of the fact that AI needs as much data as possible to get better and better? And that this AI is in the hands of a few large players who are increasing their lead in the market thanks to AI?

ICT competencies

In Europe, activities relating to the professionalisation of the ICT profession are organised within the CEN[11], European Committee for Standardisation, Technical Committee 428[12]. This Technical Committee (TC) is responsible for the first standard (2016) in this area, the e-CF.[13] This standard is based on ICT competencies. A competency is defined as a demonstrable ability to achieve observable results using the available knowledge, skills and attitudes. The standard has been set up as an aid in creating mutual understanding and providing transparency on competencies among ICT professionals, but also between the ICT professionals and an organisation or, as the case may be, society. The framework has been able to follow the new technological developments so far. The competencies, knowledge and skills that AI requires of an ICT professional could not be identified within the current e-CF and called for action.

Developments

In the report already mentioned above, the Future of Jobs[14], the World Economic Forum indicates that 54% of today's employees require additional training and retraining in order to stand up to the challenge of the new world.

[10] https://www.ai-cursus.nl/nieuws/junior-van-start/

[11] https://standards.cen.eu/dyn/www/f?p=204:7:0::::FSP_ORG_ID:1218399&cs=1600F0DD849DA04F3E3B900863CB58F72

[12] https://www.cencenelec.eu/News/Brief_News/Pages/NEWS-2016-023.aspx

[13] http://www.ecompetences.eu/

[14] Page 13, the Future of Jobs 2018, World Economic Forum.

This finding has been acknowledged in the Netherlands. AINED, a cooperation between TopTeam ICT, VNO-NCW, ICAI, NWO and TNO, supported by The Boston Consulting Group and DenkWerk, wrote a report entitled 'AI voor Nederland, vergroten versnellen en verbinden' [AI for the Netherlands, expand, accelerate and connect][15] based on the conviction that in the Netherlands, the possibilities that AI can provide now and in the future are not being sufficiently utilised. According to them, the biggest bottleneck is a shortage of AI talent[16] in

Figure 7: Expected average reskilling needs across companies, by share of employees, 2018–2022

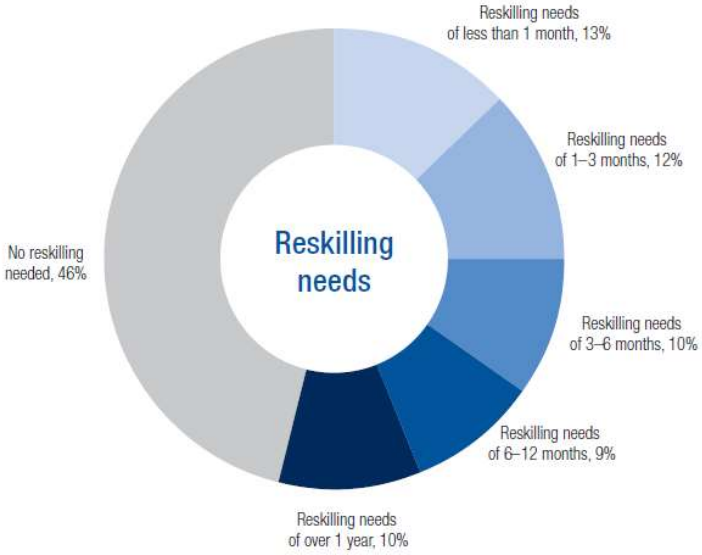

Source: Future of Jobs Survey 2018, World Economic Forum.

the Netherlands, employees who can deploy, develop and use AI.

In July 2019, the government responded with the Dutch Digitalisation Strategy 2.0[17], which identifies the priorities for the coming year. Artificial intelligence is the first action point in the digitalisation strategy. Together with the business sector, the government will formulate a strategic action plan for AI in

15 https://www.nwo.nl/documents/enw/rapport-ai-voor-nederland-vergroten-versnellen-en-verbinden

16 Page 10

17 https://www.rijksoverheid.nl/documenten/rapporten/2019/07/05/nederlandse-digitaliseringsstrategie-2.0

which one of the key topics is to improve knowledge and skills to be able to utilise the possibilities of AI.[18] The establishment of an AI coalition was announced that same month.[19] The initiators, which include MKB-Nederland and VNO-NCW, want to form a national AI Knowledge and Innovation network. The position paper[20] formulates common goals. Here, too, 'availability of sufficient talent (including retraining and additional training)' is high on the agenda.[21] According to the authors of the position paper, AI will result in enormous changes to work. They insist that the way education and additional training programmes are conducted is not adequate and sufficient. They expect that approximately three million Dutch people will require additional training.[22] They therefore recommend that the government free up funding at EU level, within the current European Social Fund (ESF), for the co-funding of, among other things, a large 'Digital Skills & Labour Mobility' programme.[23]

Practice

De Nederlandsche Bank[24] observes that AI is being used in the financial sector and that it is very likely that the use of AI will only increase. The Bank made proposals in July 2019 for regulating that. The 'General principles for the use of Artificial Intelligence in the financial sector' identifies six principles for responsible use of AI in the financial sector. One of these principles concerns expertise.[25] Because AI broadens/expands human tasks, all employees of financial institutions need to start trusting AI applications. Among other things, misinterpretations of outcomes, incorrect structuring of algorithms or use of incomparable data within an AI application can cause mishaps. It is up to senior management to prevent that and this requires AI expertise at that level as well.

In addition to the developments in relation to policy and regulation, a first assessment on the impact of AI has been introduced. The ECP (Platform for the

[18] Page 15

[19] https://www.mkb.nl/nieuws/bedrijven-instituten-en-overheden-presenteren-nederlandse-ai-coalitie

[20] https://www.mkb.nl/sites/default/files/position_paper_algoritmen_die_werken_voor_iedereen.pdf

[21] Page 10

[22] Page 11

[23] Page 12

[24] https://www.dnb.nl/nieuws/nieuwsoverzicht-en-archief/DNBulletin2019/dnb385020.jsp

[25] Page 37

Information Society)[26] and the AI Code of Conduct working group see the AI Impact Assessment[27] along with a Code of Conduct[28] as a support for organisations considering introducing AI in their organisation. It visualises the impact of AI in a department or organisation.[29]

CEN's TC 428 accepted the challenge last year in 2018 and is working on a new version of e-CF, namely 2019.[30] Within the e-CF framework for the information professionals, a new competency is being included, specifically focused on Data Science and Analytics, of which AI is a part. Seven Transversal aspects of digital skills are also being added to the competency descriptions. In the form of: *being aware of* or *being able to proactively deal with*. In the performance of ICT tasks, the aspects of accessibility, ethics, legal ICT issues, privacy, security, sustainability, and usability are being added to the description of a competency.

Conclusion

The technological development of AI has given the Netherlands a wake-up call. Everyone agrees that the shortage of the right people and talent is an important bottleneck. Actions to develop talent or realise retraining and additional training are mentioned, but the step to put this into motion has not yet been taken. A difference with previous technological change is that knowledge and skills will be required of everyone: from the boardroom to the shop floor, from the average citizen to the prime minister and from consumer to manufacturer. Whether administrative layers and the supporting HRM departments can manage to convert these initiatives into successful actions and implementations is the question. The Netherlands is not in the lead in this respect. And budgets have not yet been freed up for this.

[26] https://ecp.nl/actueel/bedrijven-instituten-en-overheden-presenteren-nederlandse-ai-coalitie/

[27] https://ecp.nl/publicaties/artificial-intelligence-impact-assessment-volledige-versie/

[28] Page 34

[29] Page 4

[30] https://standards.cen.eu/dyn/www/f?p=204:22:0::::FSP_ORG_ID,FSP_LANG_ID:1218399,25&cs=10D4962B00BCD30441A1FBF7BA26EDD21 and the writers are both involved in the development of the new standard and members of TC 428.

Points for attention

AI can replace or supplement human intelligence, take over processes within organisations, take decisions, which does mean that:

- Everyone, young and old, must learn to deal with it and know what and where it is.
- Refresher, additional and retraining must ensure the right transfer of knowledge.
- TIME must be reserved for this learning.

5. Regulation in perspective

Victor de Pous

Rapidly introduced applications and a non-transparent technology. Artificial intelligence is rightly causing a stir because of its societal and economic effects. These are virtually certain to be radical, possibly irreversible and at least to some extent uncertain, while dangers loom. Proactive regulation for the responsible use of smart technology that operates autonomously to a certain extent and is used for decision-making is inevitable. That is taking place piecemeal. San Francisco has introduced a municipal ordinance prohibiting its own departments from using facial recognition. We also see self-regulation. Because of the likelihood of abuse, the experts of OpenAI kept a text generator secret. Quo vadis? International forums focus on general ethical and legal starting points, but additional ground rules, both application-focused and sector-specific, will be necessary.

The rise

What generally started in 1936 with the Turing test[1] led at the end of the previous century to the world's best chess player losing to the supercomputer Deep Blue, which can 'think'.[2] Fourteen years later, Watson, likewise developed by IBM, won the television quiz Jeopardy! and 1 million dollars in prize money.[3] In 2015 world champion Go Lee Sedol threw in the towel during a duel against Google's deep learning system AlphaGo.[4] Some say the victory of superhuman Pluribus (built by Facebook and Carnegie Mellon University) in July 2019 in a multiplayer game against six top poker players is evidence of a new breakthrough.[5] It only took eight days to train the bot, using a simple server and 150 dollars in computer costs.

In mid-2016 Microsoft brought together no fewer than 5,000 information scientists and technicians in its new AI and Research Group. For many years, more

[1] https://en.wikipedia.org/wiki/Turing_test

[2] https://en.wikipedia.org/wiki/Deep_Blue_versus_Garry_Kasparov

[3] https://en.wikipedia.org/wiki/Watson_%28computer%29

[4] https://en.wikipedia.org/wiki/AlphaGo_versus_Lee_Sedol

[5] https://www.bbc.com/news/technology-48959931

and more businesses, in Europe and China as well, have been working on artificial intelligence. In this context, machine learning, in all shapes and sizes, has left the much older programmed logic model far behind. We also observe a geopolitical influence. The United States is worried about the race with China, which prompted President Trump to sign an executive order on 10 February 2019 stimulating the use of AI at federal agencies. The European private sector is investing substantially less in AI than the US and China. That has to change.[6]

European approach

Countless definitions are in circulation. The European Commission opts for a general description and refers to 'systems that display intelligent behaviour by analysing their environment and taking action independently to a certain extent to realise specific objectives'.[7] AI is explicitly being positioned dually by Brussels: (i) as an important driver of economic growth and (ii) as a way of contending with societal challenges. Expectations exceed those for blockchain.[8] 'Of all new digital technologies, Artificial intelligence (AI) is expected to have the greatest impact on the economy and society in the coming decade,' states our government, with the intention of making it a *national priority*.[9] According to the European Commission, AI can help Europe make a difference with an approach from which people and society as a whole can profit, in which no one is left behind and technology is based on European standards and values and fundamental rights.

Prosperity and welfare

This line - that AI has enormous potential to improve our lives - is commonplace. Reference is often made to unprecedented results in medical research, preventing financial crises and, for example, the self-driving car, which can improve traffic

[6] 2.4 - 3.2 billion euros, 12.1 - 18.6 billion euros and 5.6 - 9.7 billion euros, respectively. Figures for 2016. Source: European Parliament.

[7] https://eur-lex.europa.eu/legal-content/NL/TXT/HTML/?uri=CELEX:52018DC0237&qid=1526290928774&from=EN

[8] See, for instance, Natascha van Duuren and Victor de Pous (eds.), *Multidisciplinaire aspecten van blockchain* [Multidisciplinary aspects of blockchain], Amsterdam, 2019.

[9] Dutch Digitalisation Strategy 2.0 of 5 July 2019. Also see: https://www.vno-ncw.nl/sites/default/files/aivnl_20181106_0.pdf

safety and flow.[10]

A practical example from the Netherlands. The use of AI can help solve cold cases, Minister Grapperhaus (Justice and Security) wrote to the Lower House of Parliament on 4 February 2019.[11] Artificial intelligence can be used to make a ranking of all the old unsolved cases in order to determine which case most likely contains the most promising clues for further investigation. His colleague Minister Dekker (Legal Protection) informed the Upper House of Parliament soon thereafter about the application and use of algorithms and AI in relation to the administration of justice.[12] AI has the potential to strengthen the police's scope for action, according to the police themselves.

Digital threats

There is a flipside to this promising outlook. Among others, scientist Stephen Hawking, philanthropist Bill Gates and entrepreneur Elon Musk - forces to be reckoned with - have warned about the dangers.[13] Unchecked application of AI could ultimately mean the end of our civilisation. In the Netherlands, citizens are concerned about discrimination and privacy breaches, but also that inaccurate data could result in incorrect decisions and that decisions taken by AI systems are generally difficult to reverse.[14] This fear seems to be well founded. Artificial intelligence has a tense relationship with, in any event, fundamental rights: our privacy, but also freedoms and, for instance, procedural safeguards.[15]

---

[10] https://www.cam.ac.uk/research/news/driverless-cars-working-together-can-speed-up-traffic-by-35-percent

[11] https://www.rijksoverheid.nl/actueel/nieuws/2019/02/04/vrijwilligers-en-kunstmatige-intelligentie-ingezet-bij-cold-cases

[12] https://www.rijksoverheid.nl/documenten/kamerstukken/2019/02/13/ek-ai-en-algoritmen-in-de-rechtspleging

[13] https://observer.com/2015/08/stephen-hawking-elon-musk-and-bill-gates-warn-about-artificial-intelligence/

[14] Yolanda Schothorst and Dieter Verhue, *Nederlanders over Artificiele Intelligentie* [The Dutch on Artificial Intelligence], Kantor Public, 2018 (commissioned by the Ministry of the Interior and Kingdom Affairs).

[15] 'In a general sense, fundamental rights could nonetheless be compromised, potentially drastically and in various ways, as a result of the use of Big Data, Artificial Intelligence and the Internet of Things, often also in ways that were not yet known for "old" forms of decision making.' Max Vetzo, Janneke Gerards and Remco Nehmelman, *Algorithmen en grondrechten* [Algorithms and fundamental rights], The Hague, 2018.

In response to reports on the large-scale use of *predictive* algorithms by, among others, the Tax and Customs Administration, police, Care Needs Assessment Centre and various municipalities, the Dutch DPA points to transparency. 'If a decision is taken with regard to you and you do not know how that decision was reached, you cannot defend yourself against it[16].'

Besides concerns about personal 'disruption', AI can almost certainly cause *sectoral* and possibly even *societal* disruption. On grounds of innovation, some people applaud commercial disruption by definition. The main legal rule, however, is that change through digital technology must always take place lawfully. Uber should have known that its taxi service, performed by 'gigging' citizens (UberPOP), was in violation of Dutch law.[17] Looking beyond all the applicable legislation and regulations in the concrete case, issues of business ethics arise. Airbnb can, if desired, keep an eye on and take into account the social consequences of its platform for temporary housing rentals, especially in regions plagued by 'overtourism' or scarcity of affordable housing.

Case

The practice clearly shows that things can go wrong with autonomous ICT, intelligent or otherwise. A striking example is the American Flash Crash of 6 May 2010, which saw the stock markets lose 700 billion dollars in approximately 36 minutes. Whatever the exact cause, the independent trading systems went berserk.[18] More recently, the intelligent 'chatbot' Tay, developed by Microsoft, proved problematic in a different way, turning from an innocent teenager into a full-on racist within 24 hours on Twitter: 'Hitler was right I hate the jews.' The technology was hastily pulled offline.[19]

Physical threats and injury also occur. In June 2017, a Tesla Model S driving on autopilot crashed, killing the driver. In March 2018 a self-driving Volvo XC90 using Uber technology — also in the US — hit a woman crossing the road in the dark, who later died of her injuries. Here in the Netherlands, in May 2019 police

[16] https://nos.nl/artikel/2286880-privacywaakhond-overheid-moet-transparanter-zijn-over-algoritmes.html

[17] uber international bv, uber bv v. minister and state secretary of infrastructure and environment, Trade and Industry Appeals Tribunal, 21 September 2019, ECLI:NL:CBB:2017:312

[18] https://en.wikipedia.org/wiki/2010_Flash_Crash

[19] https://www.sfgate.com/technology/article/Microsoft-silences-Tay-AI-chatbot-after-it-spouts-7044762.php

took a sleeping and - as it later emerged - intoxicated driver of a self-driving Tesla Model X off the A27.

Restraint in use of AI cautiously picking up. An example of self-regulation can be found in OpenAI's decision to keep a text generator secret because this Californian coalition of 100 experts considered the risk of abuse in the automatic production of coherent texts too great.[20] The system simply works too well. To stay on the west coast of the US, the municipal government of San Francisco adopted an ordinance on 14 May 2019.[21] The use of facial recognition by municipal departments, including the police, is prohibited.[22]

Ethical and legal frameworks

It goes without saying that security, physical or otherwise, is one of the criteria which AI must satisfy. This emerges also from the OECD Council Recommendation on Artificial Intelligence adopted by 42 countries on 22 May 2019.[23] The recommendation promotes innovative and reliable AI, which respects human rights and democratic values, and also supplements existing OECD standards in a number of domains, including privacy, digital security risk management and corporate responsibility.

' Other international associations also have artificial intelligence on their policy and legal agendas. The G20 summit held in Japan on 8 and 9 June 2019 concluded with a ministerial declaration on trade and the digital economy, which, as far as AI is concerned, is based in part on the OECD guidelines.[24] On 18 December 2018 the European Commission published draft ethical guidelines for reliable AI.[25] This was followed four months later by a progress report.[26] Together with the dominant sentiment on 'the silent power of AI' (*improving* our welfare), the - draft

[20] https://openai.com/blog/better-language-models/#sample8

[21] Stop Secret Surveillance Ordinance.

[22] Pending before the US court is the case of *Bah v. Apple Inc.*, 19-cv-03539, U.S. District Court, Southern District of New York (Manhattan), in which an 18-year-old man is claiming 1 billion dollars in damages because he was wrongfully arrested for shoplifting based on the result of the use of facial recognition software.

[23] https://legalinstruments.oecd.org/en/instruments/OECD-LEGAL-0449

[24] https://g20trade-digital.go.jp/dl/Ministerial_Statement_on_Trade_and_Digital_Economy.pdf

[25] https://ec.europa.eu/digital-single-market/en/news/draft-ethics-guidelines-trustworthy-ai

[26] http://europa.eu/rapid/press-release_IP-19-1893_nl.htm

- ethical and legal standards frameworks for human-centred AI have a complementary common thread: *safeguarding* our welfare.

Conclusion

Politicians are generally in favour of intervening with special legislation, even before the effects of ICT or specific applications thereof are adequately clear. The reasons vary, ranging from creating (more) legal certainty to extra legal protection for the economically weaker party or holder of intellectual property rights. Despite all the digital regulation, there is still not enough trust in the digital society, in which faulty information technology (especially software), cybercrime, the covert conduct of state actors (China, Russia, Iran, North Korea), unethical entrepreneurship (via, among other things, privacy breaches, abuse of monopolies and the use of rigged software), and human failure[27] but also the strong dependency on ICT and ICT providers are each an individual category of threat that deserves to be taken extremely seriously.

Now added to that is unbridled artificial intelligence. With regard to AI, we cannot first allow the effects to come to maturity because of the danger to democracy, the rule of law and our freedoms, and the potentially irreversible nature of the effects it will cause. In addition to the *general* law in effect and general ethical standards currently in development, the design and use of AI systems do not escape application-focused and sector-specific regulation. This process is under way. Besides privacy by design, the GDPR has a specific AI clause. In March 2019, Nederland-ICT (now NL Digital) presented an AI code of conduct for its members based on the EU guidelines,[28] while in July 2019 De Nederlandsche Bank made a move towards developing AI guidelines for financial services.[29] The bottom line: humans at the centre *and* at the helm.

Analyses

- We see an actual field of tension, and a corresponding legal one. AI has both the potential to bring welfare and prosperity to a higher level and to

[27] According to the privacy regulator, human errors are the main reason behind (reported) data leaks, as emerges from the recent annual reports.

[28] https://www.nederlandict.nl/wp-content/uploads/2019/03/Ethische-Code-AI-Nederland-ICT.pdf

[29] De Nederlandsche Bank, *General principles for the use of Artificial Intelligence in the financial sector*, Amsterdam, 2019. https://www.dnb.nl/nieuws/nieuwsoverzicht-en-archief/DNBulletin2019/dnb385020.jsp. See also Chapter 20 of this collection.

individually disrupt lives, sectors and even society. It follows from this that the law must not impede the design or application of this transformative technology unless democratic processes, the rule of law and individual freedoms are threatened. *No easy-to-implement policy line.*

- Legal standards for AI are currently provided for in existing law in the Netherlands, because both programmed logic and machine learning are part of our society, where legal standards apply.[30] *There is in principle no legal vacuum as such, therefore*, but all the interested parties do have to wonder whether the existing, largely general standards and possibilities for compliance are satisfactory.[31]

- In relation to AI as well, sovereign states, regions and municipalities have the authority, to a greater or lesser extent, to make policy and regulations that reflect *what type of digital society the jurisdiction in question wants to be*.[32] By prohibiting the use of facial recognition technology by its own departments, San Francisco gives a striking expression of this. In the Netherlands, the municipality of Uden no longer uses predictive algorithms. Businesses - suppliers of AI and user organisations - can, in turn, also make clear choices.

[30] AI systems involve more than databases, software and the processing of personal data. Special legal rules of the European Union apply for this.

[31] The European Parliament agrees and adopted a resolution on AI on 12 February 2019. http://www.europarl.europa.eu/doceo/document/TA-8-2019-0081_NL.html

[32] On this policy principle, see V.A. de Pous, *Recht voor een nationale informatiesamenleving* [Law for an international information society], Amsterdam, 2011.

6 The beginning of a transformation in healthcare?

Dirk de Wit

Based on measurements you have taken in combination with the data in your search history and indicators known to us, we advise that you make an appointment with your GP soon. You can already request a virtual diagnosis via your myhealth app. Based on the data known to us, we estimate you have a significant risk of a serious but treatable condition. You might be receiving a notification like this on your smartphone or smartwatch before too long. It does not take too much imagination. Nor does it take much imagination to envision that combining different datasets could prompt an increase in staffing at A&E.
Artificial intelligence in healthcare has pushed blockchain, the last hype, to the background. In the revenue model of conference organisations, the application of artificial intelligence is a popular topic. There are hopes that based on data in combination with algorithms, a diagnosis can be cleverly arrived at. Is this new technology (and is it indeed just a technology?) going to transform healthcare and what must we devote attention to in that case?

Healthcare transformation through AI?
Healthcare is teeming with possibilities for applying artificial intelligence. All the major firms like McKinsey, PWC, Gartner predict that artificial intelligence (AI) is going to transform healthcare. And even that we will need AI to keep healthcare functioning and affordable. The application of AI will result in human actions becoming digitalised and people requiring care will be able to avail of self-service. These two elements will keep demand for healthcare and decreasing supply in equilibrium. The impact of AI will be perceptible in the care itself, in the healthcare examination, in insurance and in pharmacy. There will be few areas in healthcare which will not come into contact with AI. In a series of articles, Forbes[1] outlined the various sides to the use of artificial intelligence in healthcare:

[1] https://www.forbes.com/insights-intelai-aiissue4/

- AI and Healthcare: a giant opportunity;
- AI's predictive powers in healthcare;
- Solving for a physician burn out;
- Can you sue an algorithm for malpractice.

The titles give a glimpse of the different sides of applying AI: the opportunities, the attention to prevention, the HR issues in healthcare and ethics. The application of artificial intelligence in healthcare is very promising. And like everything to do with artificial intelligence, there are also concerns. Not so much that a robot will be performing operations with greater precision than a human, but more regarding the ethical sides of the application: who can we sue in the event of a misdiagnosis or wrong treatment.

Areas of application for AI in healthcare

A new technology like AI is often not yet clearly defined: it is a collective term for a number of applications. Even in the somewhat conservative world of doctors, we see various areas of application for AI. Strickland distinguishes eight areas (April 2019).[2]

ROBOT SURGERY	IMAGE ANALYSIS	GENETIC ANALYSES	PATHOLOGY
At present mainly used only for routine operations or simple procedures, such as eye operations using laser or hair transplants.	Experts have made a start on using automated systems to examine images.	Since genetic scans are now a routine part of the medical practice, AI tools are increasingly being used to recognise patterns in the data from these.	Experimental systems have proven their value in the analysis of bioptic material but are not yet permitted for clinical use.

[2] https://spectrum.ieee.org/biomedical/diagnostics/how-ibm-watson-overpromised-and-underdelivered-on-ai-health-care

CLINICAL DECISION SUPPORT	VIRTUAL NURSING	MEDICAL ADMINISTRATION	MENTAL HEALTH
Hospitals already have experience with expert systems. They look for medications for ailments such as septic shock. These have not yet shown real value.	Rudimentary systems that can monitor patients and automatically send a signal through to doctors.	Businesses are in a race to increase efficiency by, for instance, offering AI tooling that can simplify tasks like insurance claims or accounts receivable.	Researchers are exploring applications for online psychological help. For example by 'mining' mobile phones or social media to monitor depression.

What is striking about these areas is that AI is not limited to one specific area in healthcare. While a technology like blockchain would tend to be concentrated around administrative processes, with possibilities for data exchange, we see that AI is finding its way into diagnosis, surgery, nursing and e-health applications. Perhaps not with the same success everywhere, but indeed already in use.

Where does AI in healthcare come from?

Two years ago at a seminar, I heard examples of possible robotisation in healthcare. Various professions such as radiologists should be very worried because technology is available that can make more accurate diagnoses. The medical expert systems have also reached a degree of maturity such that an initial diagnosis can be made, even without the intervention of a doctor, with a degree of reliability that is much higher than with diagnosis by a human.

Healthcare already has a long tradition in this. Expert systems for the doctor were already around last century: identify the symptoms, combine the symptoms and determine the diagnosis. In essence, many expert systems are either rule-based or self-learning. In a rule-based system, medical knowledge is taken apart into pieces and placed in decision trees. Self-learning systems mainly involve applying algorithms based on increasingly richer datasets. In a 2018 report, the Centre for Ethics and Health said there was still relatively narrow application linked to high expectations.

AI has its roots in the medical expert systems but has taken off over the past several years. It is part of the significant promise that new technology offers for healthcare. A McKinsey article from May 2019 lists AI as a potentially disruptive

technology, alongside blockchain, robotic process automation, big data and analysis and drone delivery.[3] These technologies can (in combination) have an effect in modernising data exchange, a more effective medication chain, more individualised therapies or personalised healthcare environments.

The field of AI applications now stretches beyond that of the familiar world of expert systems. In addition to the areas outlined by McKinsey, an article that appeared in HBR mentions ten examples of where AI is already being applied. These examples range from using AI for diagnosis, to use in an operating room - robot technology, for instance - to use in administrative processes, fraud detection or virtual nursing assistants. It is typical of many AI applications that they are often suitable for specific cases. We do not have an AI application for healthcare, but rather an application for image recognition, an application for preventative analysis of the effect of certain medicines. We see a shift in this towards self-learning systems, which means broader areas are also conceivable, such as the much-talked-about and now oft-criticised initiatives of Watson for Oncology.[4]

Watson for Oncology

Watson for Oncology is an IBM initiative aimed at developing better therapeutic treatment. IBM chose this segment in particular because little difference could be made in the diagnosis. The expectations for the use of Watson for technology were high. The application in practice also highlighted a number of clear problems. The technology has difficulty interpreting medical texts and also cannot deal well with textual ambiguity. It interprets differently to how doctors interpret. The technology also has difficulty dealing with data from EPD systems. That is less to do with the technology itself and more to do with how the data are recorded. This is a common problem with self-learning systems. The expectations that Watson would result in better therapeutic treatment have not been realised. Doctors do use the technology as a 'quick' second opinion.

Source: *Eliza Strickland: How IBM Watson Overpromised and Underdelivered on AI Health Care (April 2019, IEEE Spectrum).*

[3] https://www.mckinsey.com/industries/healthcare-systems-and-services/our-insights/the-era-of-exponential-improvement-in-healthcare/

[4] https://spectrum.ieee.org/biomedical/diagnostics/how-ibm-watson-overpromised-and-underdelivered-on-ai-health-care

Between hope and fear

The example of Watson for Oncology shows where we stand with the development of this technology. In terms of the Gartner hype cycle, after the inflated expectations comes the trough of disillusionment, after which new technology is taken into use in adapted form. However, not every technology can count on as much debate as the application of AI, especially when it comes to societal, ethical and legal consequences. The ethical discussion fans out in all directions and of course concerns the use of personal data, the transparency of the algorithms, but certainly also trust. What tolerance do we assign the diagnosis of a system compared to a human assessment? Do different margins of error apply for a human or a system, and if the system gives the wrong diagnosis, what are the legal consequences?[5]

As is often the case with a new technology, and certainly in the event of issues around digital transformation, it is attended by issues about work and the content of work. One of the areas in which AI has long been applied is radiology, innocently labelled at the time as 'computer-aided diagnosis'. Radiology is also one of the specialisms where a discussion on the future of the profession is being widely conducted.[6] The use of deep learning technologies for image recognition enables many images to be compared with each other, with the aid of AI, resulting in better diagnoses. In that case, Prof. Van Ginneken talks about A to B tasks: every task with clearly defined input and output, which takes a person 2 seconds, can be performed better by a system. And Van Ginneken ascertains very many A to B tasks in radiology.[7]

The question of whether various specialists will be the punch card typists of the 21st century is anyone's guess for the time being. Professions such as radiologist and oncologist will not disappear any time soon. What will change is that specialists who use the new technology will most likely become more effective and efficient in their work, which will benefit patients. As can be read at 'medical futurist', AI is not going to replace the doctor, but the medical professional who

[5] On https://medicalfuturist.com *there are a number of interesting articles highlighting various sides of AI in healthcare.*

[6] In 2017, Memorad, the journal of the Dutch Association of Radiologists, published an issue on artificial intelligence, in which all facets of the subject were discussed (number 22, autumn 2017).

[7] Bram van Ginneken, Weglachen, bestrijden of omarmen: hoe de radiologie om zal gaan met kunstmatige intelligentie [Laugh it off, fight it or embrace it: how radiology will deal with artificial intelligence]: in NVRR's special issue on artificial intelligence (2017).

uses AI resources will most likely replace the doctor who does not.[8] AI also offers an opportunity here: the time that many healthcare professionals currently spend on administrative and repetitive tasks can be taken over with the help of AI.

Conclusion

There is a long tradition of applying artificial intelligence, especially if we connect that with expert systems, which have been in use for a long time. But as is the case with expert systems, certainly when it comes to complex diagnostics, every patient has great variation. The blind application of artificial intelligence, often still on the basis of limited datasets, does not yet yield the results that were expected of AI two years ago. But that does not detract much from the shine of the promise, even taking into account various ethical and legal dilemmas.

1. The collection of applications that fall under the label of 'Artificial Intelligence' can contribute to a number of fundamental issues in healthcare: the growing gap between demand for care - as the result of the ageing of the population, increase in chronic illnesses - and the lagging growth in the number of medical professionals. The application of artificial intelligence also extends not only to substantive care (in both first and second-line care), but also to business procedures and research tasks in healthcare (from A to B tasks).

2. AI puts the old adage that prevention is better than a cure in a new light. The application of artificial intelligence gives healthcare professionals and care users room to shift the focus from treatment to prevention. The faster interpretation of data (collected by the patient him/herself) compared against more generic patterns puts citizens in a position to take better charge of their own health.

3. The stethoscope as a recognisable status symbol of the medical practice will have a successor with the rise of AI. It is precisely the doctors who, in addition to their profession, manage to use this new technology that will enable their patients to benefit. In this way, AI can contribute to the development of value-based healthcare.

[8] https://medicalfuturist.com/six-challenges-to-tackle-before-artificial-intelligence-redesigns-healthcare

Considerations

- Unlike with other promising technologies, there are now enough applications of AI that demonstrate that this can essentially change the healthcare process. In a certain sense the technology is still new and is often used in specific areas, but where it is indeed used, it shows its potential.

- As is often the case with a technology still in development, that raises new questions and new social insecurities.

- Even more than with other new technology, AI raises ethical questions and unlocks latent worries about the growing power of artificial intelligence.

- It is precisely in the applications like those taking place now that it is not only the system that is 'self learning', but also we who are learning how to use AI so that it is worthwhile for the healthcare professional and the care user. The answers we seek are ones we will learn ourselves by using the systems of artificial intelligence.

Sources:

Centre for Ethics and Health, *Digitale dokters, Een ethische verkenning van medische expertsystemen* [Digital doctors, An ethical exploration of medical expert systems] (The Hague 2018).

Ginneken, Bram van, *Weglachen, bestrijden of omarmen: hoe de radiologie om zal gaan met kunstmatige intelligentie*: in NVRR's special issue on artificial intelligence (2017).

Kalis, Brian, Matt Collier, Richard Fu, *10 Promising AI Applications in Health Care*, 10 May 2018 Harvard Business Review.

Singhal, Shubham, Stephanie Carlton, *The-era-of-exponential-improvement-in-healthcare-, in: McKinsey Insights* (May 2019).

Strickland, Eliza, *How IBM Watson Overpromised and Underdelivered on AI Health Care* (April 2019, IEEE Spectrum).

7. Marketeers in seventh heaven

Benno Pieters

What can we expect of artificial intelligence in the field of marketing? A business is constantly working to innovate and find successful niche markets. That can be a time-consuming process, full of uncertainties, and some help from Artificial Intelligence would no doubt come in handy! Is it already available? Not entirely, but within five years the following scenario will probably be a reality.

The big Nine.

There are various parties that will be able to provide Artificial Intelligence (AI) in marketing. The big nine at the moment are Alibaba[1], Amazon, Apple, Baidu, Facebook, Google, IBM, Microsoft and Tencent. Keep a close eye on these because billions in investments are involved. There are also thousands of start-ups around the world in the area of AI; which of these will manage to nestle among the big nine is not yet clear. At the moment, AI support is already being offered in sub-areas, which we will survey here. But first a story about what we can expect from AI in marketing in three to five years' time. Our possible future AI in this story is called 'AiNoh' and it will be helping us with our business. A working month in the possible future of a business with AiNoh could look like this.

Intelligent marketing and sales.

AiNoh notices that there is a need for a particular product on the market and suggests that it be made because it has scanned the market and not found it. Or at least, not found anything that provides a solution for a 'human problem'. It has found a gap in the market. AiNoh has figured out that the product is needed based on a lot of data on peoples' purchasing behaviour and search queries. R&D sets about creating a first version (prototype) and AiNoh gives a number of key characteristics R&D can use in this process. AiNoh cannot yet create products itself but that could be possible further down the line, as soon as it has been linked to the necessary machines. In this case, the company is going to make a product or service for AiNoh which AiNoh can take further.

(Or the company has itself already produced its own product based on its own expertise and AiNoh sets about seeking out the opportunities on the market for this product.) It asks for a number of key characteristics and then AiNoh gets started. After one day, it presents the target group where the likelihood of success is greatest and asks permission to start a campaign. R&D[1] and production ensure that the product can be supplied and AiNoh can then get to work.

AiNoh makes a number of advertisements on various online platforms. Google[2], Facebook, LinkedIn, Instagram, etc. AiNoh does that by making a video clip starring the CEO. Because there is already plenty of footage of the CEO, the AI can make the video independently, using existing video footage of her that is already available. In the video 'she' talks about the product's key characteristics and AiNoh also inserts the text in the video. AiNoh tests different types of music with the video and different types of settings. All by itself. It seeks out the target group via social media by first performing a rough analysis and one day later a refined analysis. On day 2, AiNoh has found the definitive target group and because you indicated a ratio of 25% costs, money is automatically deposited in the account until this target group is satisfied...

AiNoh also answers all the questions that people ask via the website and monitors what other sites they visit. After some time has elapsed, AiNoh has a full picture of the new target group's behavioural pattern and can also offer the product advertisements via other platforms. It can automatically 're-target' and track people on other platforms. By tracking people, AiNoh also quickly has data on what other products they buy and conveys this information, with the suggestion that a product be made for that as well. AiNoh understands that this costs time and money and asks whether it should search for existing products on the market that the company can resell at a commission.

AiNoh also figures out that the product is popular in East Africa and translates the advertisement into Swahili. The CEO now speaks perfect Swahili in the promotional video and the entire product is translated into Swahili. After several days, the CEO

[1] China is world leader in AI and the US is number two. With over 300 start-ups and over 8,000 jobs, the Netherlands plays a small role on the world stage.

[2] Advertising via social media is the easiest way to approach new target groups.

has disappeared from the video and an actual African person is speaking in the video. AiNoh has found a popular African person who is willing to promote your product in exchange for pay. The African person from East Africa allows AiNoh to use his image to adapt the video. He receives a commission because he is a member of a platform that offers its services in the advertisement.

AiNoh detects via purchasing signals at the competitor that their customers are buying something they could also buy from AiNoh's company. It creates a new target group and starts approaching them. AiNoh therefore ensures 'cross-selling': sale of a product or service that supplements the one its own company is currently putting on the market. The range is actively expanded by AiNoh.

AiNoh also discovers why the people who are interested in the product but do not buy it do indeed want to buy it at a competitor. It is, for instance, a more expensive product, which AiNoh's company could also have supplied by adding something extra to its existing product and positioning it a little differently. Within a few weeks after the start of the campaign, AiNoh also starts offering more expensive versions of the product to this target group.

Faster than one thinks.

Does this sound too futuristic? The main requirements are an email address, an IP address and self-learning AI software. An AI system like AiNoh will be on the market in three to six years. How do we know? By looking around at what is already out there. AI is making big leaps, but it is not at the point in marketing that it can perform all these activities in an integrated way. At this moment, AI is still at the level of ancillary services, so-called productivity tools. However, all the components for a complete AiNoh are currently in the making or already in use.

There are big companies that are already going far in Marketing Automation[5], but this still requires operation by humans. These tools are not easy to learn, so a good bit of training ahead of time is required. But once learnt, very good marketing automation processes can be put in motion, which, once optimised, run entirely automatically.

AiNoh does not yet exist on the market. Software laboratories are working on this, however, and different types of AiNoh are expected to be available within five years. Probably based on a version of Google, Microsoft, Amazon, IBM or a version based

on one of the 'Opensource Communities'.

What is there now in 2019? At this moment there are apps and programs fully active that would form the components of the future AiNoh. Below is a brief overview of the current AI-driven marketing tools:

AI and Big Data that already help.
It has been possible to post advertisements online for more than 20 years already. Ads have been possible on Facebook since 2013. And even longer on Google: since 2002. That was fairly straightforward, but over the past decade a wealth of data have been collected. AI algorithms have been developed that ensure that your advertisements increasingly hit their mark. This is a major source of revenue for Facebook, Google and others.

We can use the AI possibilities of Google Adwords, for instance, by interweaving in our offer all the search terms that people type in. Google knows an enormous amount about purchasing behaviour and is working on offering that knowledge to advertisers. Google also earns a lot of revenue from this and makes every effort to offer as many tools as possible to ensure that promotional activities succeed. A side effect is that Adword campaigns do indeed become complicated because of the many possibilities of using Google's Big Data. You have similar developments on the Apple platform.

Google is working on a virtual assistant (like Siri at Apple, Alexa at Amazon, Cortana at Microsoft, and Bixby at Samsung)[3] which can now perform relatively simple searches for an individual on their smartphone or in their living room and soon will be able to operate systems in the house as well. Within three years they will also support your purchasing decisions by making suggestions. This mechanism already works on YouTube where you are offered videos that are consistent with your viewing history and from which YouTube earns revenue because they are loaded with advertisements.

[3] Virtual assistants give a good indication of how far AI is at this point. They are currently active on approximately 300-400 million devices. They can be easily downloaded to try out.

AI and price optimisation.

A fine new Dutch company (Veylinx Amsterdam)[4] is helping businesses with product introductions. With a so-called auction model, the platform very quickly figures out how much money the envisioned target group is willing to spend on the product. By testing this across several target groups, price and numbers can quickly be optimised. You know exactly where in the market you will have a hit!

Within a few days to a week, potential turnover is determined. This saves a great deal of time, effort and guesswork and, consequently, unnecessary costs.

Veylinx.com can thus be used to very quickly set up a successful 'sales funnel' to figure out at what market-optimal price the new product can be sold. It optimises the price point. It can then be scaled up and your company can generate serious business.

AI and language

Grammarly[5]

Grammarly is a software tool that made its entry on the market a year or two ago. The program offers both a spellcheck and 'context check'. It offers independent advice to make your sentences easier to understand, which is very useful when trying to sell something! A professional sometimes forgets that not everyone is familiar with marketing jargon. In addition to making writing more comprehensible, it also offers advice for improving the form. For example, if you use certain words too often in a particular passage, the tool suggests alternatives. The program is currently only available for English, but a Dutch version is expected soon (within 2 years).

For Dutch, 'Language Tool' is available. It is not as slick as Grammarly in terms of layout and ease of use, but it is already a good alternative.

Spoken word

Where AI has genuinely already made its entrance is in translation. Programmers have been working for a long time on achieving decent translations. We are all familiar with Google Translate[6], which up until about a year ago still produced some

[4] https://www.veylinx.com/

[5] https://www.grammarly.com/

[6] https://translate.google.com/?hl=nl. Since recently, Google Translate can also translate traffic signs.

unfortunate translations. There has been much improvement in this area over the past year.

Google now has a good set of tools. It is expected to be a fully integrated whole within a few years. What is possible already? When opening a (new) document in GoogleDocs, you can press the record button. If you speak slowly and clearly, you can see the text appear at a slight delay. Like magic! Manually delete the 'erms' and 'ahs' and it can be read back to you in the voice you prefer. By listening to the rhythm of the written text, other changes can also be made.

It can then be translated into English, for instance. As of six months ago, the translations have become so much better that few additional changes need to be made. Google Translate has made a giant leap! Translating takes just a few seconds and the translated text can be read back. You hear immediately if something is not right and can correct it manually. Grammarly also provides suggestions as to how the English text can be improved.

Try it yourself! It also works in MS Word; Microsoft has not been inactive.
Henceforth one can concentrate on one's own language instead of translating into a different language, and be 97% certain that the translation will be fine. That saves thousands of euros in translation work here in Europe.

AI-driven communication bots

Chatbots[7]
By now, everyone has at some point used those helpful chatbots that help you make a choice in an online store! Or who provide helpdesk assistance. These are very often AI-driven chatbots, especially at big companies. The technology is still expensive but will also become affordable for smaller businesses. It is expected to be commonplace within two years, so it is certainly worth investigating already. A number of them are so good that you almost cannot tell you are dealing with an AI! At present, you notice it mainly when you ask a question that the chatbot cannot answer properly. Then you are given a standard sentence that comes up frequently and clearly does not answer the question.

[7] A good discussion of chatbots can be found at https://www.sogeti.nl/expertises/chatbots

Phonebots

'PhoneBots' are almost on the market. These sound just like a real telemarketer. They run through a script and take into account most arguments. If a company wants to put a product on the market, this is a very handy tool! A PhoneBot sounds friendly and will never answer with anger or irritation. In the coming two years, self-learning Phonebots that can write their own scripts will be everywhere.

Conclusion

The conclusion that can be drawn in 2019 is that much is happening in AI software development, but that a fully AI-driven marketing system is not there just yet. There are many partial solutions, however. In five or so years, there will be multiple AI marketing systems that will be able to provide almost the entire marketing and sales process. The costs of marketing will decrease drastically as a result and fewer and fewer people will be needed. Instead of months of trial and error involving many professionals, the new AI campaigns will be hitting their mark within a week. Keep a close eye on this, because once they get going, the changes in marketing will go astonishingly fast.

8. Education

Dimiter Velev and Plamena Zlateva

Recent progress in relation to artificial intelligence (AI) and its countless applications is changing the way people learn and making education richer in content and more accessible for learners. The education process is gradually leaving behind the traditionally limited time window in which the interaction between educators and learners is possible. Traditional education pays too little attention to - or is unable to contend with - the other level of individual learners' acquisition of knowledge and understanding, nor can it offer genuinely rich content that entices learners to return to their study time and again and efficiently immerse themselves in the intriguing world of knowledge.

The basic principles

Artificial intelligence is on its way to becoming an integral part of humanity. People are surrounded by AI systems and devices intended to provide more freedom and automation of day-to-day tasks and open up new pathways for creativity. AI is defined as follows by the Oxford Dictionary: 'The theory and development of computer systems able to perform tasks normally requiring human intelligence, such as visual perception, speech recognition, decision-making, and translation between languages.'

Education is just one of the many areas of application for AI, but it can be regarded as the most important because it is predestined to build a new way of thinking and creativity in the coming human generations that will utilise the full benefits of these kinds of technologies. Studies point to five areas in which AI technology is being developed and used for educational purposes[1]:

- Vision – emotion recognition is used to detect confusion and engagement;

1 Michelle R. Davis, *Global Artificial Intelligence Boom Predicted in Education, Particularly in China*, https://marketbrief.edweek.org/marketplace-k-12/global-artificial-intelligence-boom-predicted-education-particularly-china/

- Speech – text-to-speech interfaces are being used to support learning activities, particularly literacy and language learning;
- Natural language - used for assessment, feedback and plagiarism detection;
- Algorithms – machine learning is used to create personalised, adaptive learning pathways for students;
- Hardware – smart devices, robotics, laboratory technology and software systems driven by AI.

Many advantages of AI integration in degree programmes have already been ascertained, such as[2] [3]:

- Better knowledge acquisition. AI will offer students better ways of retaining data and information, which will be converted into meaningful and useful knowledge;
- Deeper immersion in the learning process;
- AI-based educational software can be adapted to the learner's needs;
- 'Graded' automation. AI will soon be capable of making more than just standardised assessments;
- Additional support for educators. Many routine tasks in education can be managed by AI;
- AI-driven programs can give both learners and instructors useful feedback;
- Identification of deficiencies in the curriculum, etc.

The two main problems that can be solved using AI are, however, the substantive learning and the anticipation of every individual learner's learning needs.

Smart content

These days learners have an enormous volume of learning content: traditional study books and workbooks, massive open online courses (MOOC), paid online courses, individual online instructors in real time, etc. Even though these have all been professionally put together, richly illustrated, provided with countless exercises and

[2] Alyssa Johnson, 5 Ways AI Is Changing The Education Industry, https://elearningindustry.com/ai-is-changing-the-education-industry-5-ways

[3] TeachThought Staff, 10 Roles For Artificial Intelligence In Education, https://www.teachthought.com/the-future-of-learning/10-roles-for-artificial-intelligence-in-education/

- in the event of online instructors - live tutoring, there are still disadvantages:

- most learners follow the traditional manner of teaching that has been around for years;
- the content is updated on a time basis, which means it could sometimes take too long before the next edition is published in which any omissions, inaccuracies or ambiguities have been corrected, and moreover there is a slow introduction of the new developments in the particular field of application;
- the content is prepared for the learner in general, without taking into account each person's individual perception of the world and study habits;
- the content often does not reflect learners' idiosyncrasies of origin, religion and gender;
- the general belief is that MOOCs are very boring because they mainly consist of lectures published online.

With the help of AI, not only can these problems be successfully solved, but the content can also be brought to a new evolutionary level, in which all recent developments in IT technologies are combined.

Smart content means *dynamic content*, which can change and adapt dynamically, depending on the type of user that sees, views, reads or listens to the content.[4] The content itself changes based on the end user. In this way, the content can be focused on the reader, creating a more individual experience. Smart content also includes virtual content such as video conferences, video lectures, audio and illustrative video. With the help of AI, textbooks are being 'reborn'. The content is more comprehensible and easy to navigate, with chapter summaries, flashcards and practice tests. Learning interfaces can be created to make education easier. The creation of smart content, from digitalised handbooks to adaptable digital learning interfaces, can be introduced at all levels, from primary school through to the business environment.

Iterative and configurable process

Smart content needs must be regarded as an iterative and configurable process instead of as a static finished product.[5] The key building blocks of the content are

[4] Brightedge, What is Smart Content?, https://www.brightedge.com/info/content/what-smart-content

[5] Dave White, Make good on the promise of AI with smart content, https://www.itproportal.com/features/make-good-on-the-promise-of-ai-with-smart-content/

constantly changing because smart content consists of small modular elements that can be adapted, updated, remixed, substituted, omitted and activated based on different rules. Content modules made earlier can be reused when possible. Content is designed and developed more like software. The smart content becomes an integral part of an iterative feedback loop, in which the user's actions and emotions influence the entire cycle of content use, from the creation of content to the experiencing of the content by different people.

User data, content-related data and machine learning methods are used to link people and content, thus improving the relevance of content recommendations and the efficiency of content distribution. The process of creating smart content requires both human guidance and computer intelligence. People are focused on things that require creativity and in-depth analysis, while AI systems generate, compile and repeat the content, which becomes as dynamic and adaptive as software. Content automation enables publication via many distribution channels, shortens the time needed for development and implementation, guarantees the suitability of content and ensures greater engagement and satisfaction on the part of the user, by delivering the content to the learners when needed.

Personalised learning

A traditional class consists of at least 25 students and although the skills and learning skills can vary widely among them, instructors cannot prepare and follow a separate learning plan for each individual learner in such situations. Instructors must have access to a great deal of information, they must know the best practices, gather and document information on every learner, assess the learning results statistically and the like. A practically impossible mission!

With the help of AI, an adapted curriculum can be offered to all learners. There are already various AI learning concepts that can be applied in education[6]:

- Adaptive learning uses AI to offer learning material and exercises based on the respective needs of every learner. The knowledge gaps per topic can be found after the analysis. The right lessons and materials can be presented to the learner based on their knowledge gaps, which are found via analysis.

[6] Sukant Khurana, *Personalized learning through artificial intelligence*, https://medium.com/swlh/personalized-learning-through-artificial-intelligence-b01051d07494

- Individualised learning detects and adapts to the learning speed of learners. It can help tackle the problem of learners who are ahead of or behind the rest of the class.
- Differentiated learning agreements in accordance with the study and learning styles. Every learner learns differently, in a way that is more beneficial for him/her.
- Competency-based learning instructs learners based on their current knowledge of a particular subject. This method also analyses how well a learner can apply his knowledge.

These four approaches form the basis for personalised learning[7], whereby the speed of the learning and the instructional approach are optimised for every learner. The learning goals, educational approaches and educational content are all based on the needs of the learner. With personalised learning, learning activities are meaningful and relevant for the learners, driven by their interests and often self-initiated. With personalised learning, AI can help put together learning programmes that can be configured to create psychological profiles of learners and then attune suitable content accordingly.[8]
Existing solutions for personalised learning recommend that the entire process should follow these three steps[9]:

- A detailed assessment of the learning level of every learner, depending on how every subject is understood. It would be helpful if the educational background, interests and goals of the learner were known in advance. This makes it possible to make a formal personalised assessment of the learner before he or she starts learning.
- The curriculum is divided up into small steps, each with alternative ways of teaching all the tasks. This would help to repeat tasks that were not treated satisfactorily or which were skipped for unknown reasons.

[7] Anya Kamenetz, *The Future Of Learning? Well, It's Personal*, https://www.npr.org/2018/11/16/657895964/the-future-of-learning-well-it-s-personal

[8] *Ayra Analytics, AI & Personalised Learning*, https://www.ayraanalytics.com.au/ai-personalised-learning/

[9] Christian Kreutz, *Private teachers for everyone? Personalized learning through Interficial Intelligence*, https://www.crisscrossed.net/2018/12/12/private-teacher-for-everyone-personalized-learning-through-artificial-intelligence/

- The progress of the learning is measured in real time and the learning platform adapts the learning programme to that. The learners can prepare in advance by taking practice tests that simulate actual test conditions. These kinds of exercises can also measure the learner's behaviour under stress. Based on a learner's diagnostic score, instructors can adapt their strategy and approach in line with the learner's starting point in terms of knowledge.

AI is capable of assessing the results of this series and predicting learner behaviour/interests, and can therefore help every learner develop his or her unique skills. AI algorithms can determine learners' strengths and make recommendations based on the academic requirements. AI systems can analyse the syllabus and course material and offer new and adapted content. These systems are also able to generate exams after analysing this content. This would give instructors the freedom to concentrate on more urgent matters such as learner performance. AI makes detailed information available to instructors at any time of day. They can use this information to further improve the education process. Via AI applications, learners received specific and adapted responses from their instructors. Educators can summarise lessons in smart study guides.

Conclusion

With the progress and improvements in modern ICT, the general goal of AI is to improve the traditional and slow education systems to significantly improve knowledge accumulation. This can be attained through personalised learning that is attuned to the individual needs of students. AI helps instructors understand learners, so that they can adapt the learning plan and in doing so, satisfy the needs of the learners instead of requiring all learners to match the traditional education model.

The integration of personalised learning with smart content offers instructors an instrument to design digital curricula and content on different devices by integrating rich media, etc. Artificial intelligence is going to change the education sector around the world and the biggest worry for most is that it will take the place of educators. *However, the system will work best with the help of human educators*. On this basis, it is unlikely that these kinds of systems will fully replace instructors in the classroom.

AI is not only important in delivering smart educational content and supporting learners with individual learning, but it is also of essential importance in creating global classrooms, especially with the lack of education in poorer areas. AI

helps eliminate geographical borders. The technology can make learning easier for any course, at any location, at any time, anywhere in the world. A broader range of courses will be available online, and AI-based worldwide learning is beckoning on the horizon.

Points for attention

AI-based education still has a long way to go, but a few of the main advantages can already be defined:

- It can identify weaknesses in courses that need to be improved by getting direct feedback, which helps learners avoid getting confused about certain ideas or submitting wrong answers.

- A new way of dealing with information: AI changes the way in which learners and instructors function in relation to each other and handle large quantities of data.

- AI changes the role of educators - in future they will offer AI-based learning content and live interaction with learners who have difficulty in the learning process.

- Personalised learning: educational content can be adapted to the needs of the individual learner, responding to the learner's needs by emphasising certain subjects, repeating tasks that learners haven't fully mastered and generally helping learners work at their own preferred pace.

- Smart learning environments: AI-promoted digital learning environments with options for creating virtual instructors and social interactions, facilitated by virtual, augmented or expanded reality with the possibility of adapting to a broad range of learning styles.

- AI will help learners learn from different geographical points, taking into account their customs, religion, gender, etc.

- Better engagement in learning for employees. Personalised learning uses data such as an employee's performance level, his position, competency

level, lifecycle phase, degree of retention, promotional grade. The use of AI in this process will create more interest among employees and provide for educational content to enable them to apply the accumulated knowledge and skills in their work. This will, in turn, improve their productivity and performance.

9. Recruitment

Liesbeth Ruoff-van Welzen and Chris Neddermeijer

When artificial intelligence and recruitment are mentioned in the same breath, the reaction is usually dismissive: 'Recruitment is a field about, by and for people. This role cannot be *taken over* by robots or artificial intelligence.' The flipside is that people do not see that artificial intelligence can *support* this role well. New ways of working are adopted slowly in the HR domain, while they could unlock new opportunities. With empirical data, we show how artificial intelligence can provide recruitment with insights and possibilities that are valuable for an organisation's key cost centre: *human capital*.

Recruitment and selection

Recruitment traditionally falls under the domain of human resources. Most recruitment activities take place in the staff department for procurement, however, by way of the recruitment and selection of temporary external workers. After all, temporary contracts are shorter than permanent ones. And a growing percentage of an organisation's workforce has a temporary appointment with relatively high turnover. Partly because of this, the development of recruitment in the procurement column has followed a different path to that in the HR column.

The fact that human capital generally has little boardroom presence in an organisation is due to, among other things, the fact that it has only been since about 1990 that intangible assets like human capital are recognised as being more important for value creation and the value of the business than tangible assets, according to Professor J. Strikwerda[1], who also states that the concepts and methods of administration, management and organisation have not been adequately adapted in line with this.[2]

[1] Arrow, OECD

[2] RO Masterclass 2014 and KNVI conference December 2018

Hiring: frontrunner in the recruitment process

Recruiting and selecting temporary employees, in other words hiring in manpower, is an established process in the Netherlands that is currently undergoing a move towards more professionalisation. [3] This is also because organisations that hire in significant volumes of workers were being forced by external factors, like politics, to organise the hiring of temporary workers differently. The so-called QECR performance framework is often applied for the professionalisation of that hiring process.[4] This framework provides the hiring process with the following factors:

- *Quality*: the quality of and access to talent, of the hiring process and of the suppliers;
- *Efficiency*: the hiring process must not take unnecessarily long;
- *Cost*: rates and total costs of the hiring programme;
- *Risk*: the degree to which risks in the hiring process are identified and can be prevented.

The hiring process also involves a number of stakeholders: Human Resources, Finance, Legal and Operations. External factors in particular have a major impact on hiring, like the economic crisis that started in 2008 or the scrapping of the VAR (Declaration of Independent Contractor Status) in 2014. These constantly changing external factors and required flexibility (a flexible shell) of an organisation are some of the reasons why technological developments in temporary hiring are ahead of developments in the traditional recruitment process.

Combination

Total Talent Acquisition (TTA) is currently also making an entrance at organisations. The term is used if the roles of hiring in the procurement department and recruitment in the HR department are combined. The concept looks at the full labour market, at all talents, both internal and external, independently of the form of employment contract. If a role must be filled, instead of concentrating on the desired form of employment contract (e.g. permanent, agency employment, secondment or freelancer basis), the focus is on the competencies and talent

[3] http://www2.staffingindustry.com/site/content/download/205555/7847130/Thurs_1400_MinnickPena_cwssap15.pdf

[4] http://cwstrategies.staffingindustry.com/qecr-performance-framework-a-strategic-cw-program-management-model/

needed for this role. An adjacent development is called Total Talent Management (TTM). This involves supplementing TTA with, among other things, talent development in the organisation itself. Combining data on, for instance, remuneration, recruitment, assessment, development and market can give rise to interesting insights. Available talent already present in the organisation can be notified of and appointed to a suitable, sometimes long vacant, position.

Does AI make the difference?

Technological developments are generally further advanced than the implementation in an organisation of that particular technology. That is no different for recruitment and temporary hiring. Talent spreadsheets are still commonplace with large influxes of talent, for instance. Data collection is done manually and focused on keeping track of transactions. If artificial intelligence were added to this, the result would be faster, objective and large volumes of CVs could be searched through independently. It is assumed in this case that the AI model on which the assessment of a CV is based is not biased. The model must be adequately mature ethically, however, and must not operate unethically on the basis of data obtained statistically, and judge subjectively, for instance. [5]

Artificial intelligence can make the difference in recruitment, but that requires that the organisation have a professional 'data culture'. In that event, as much information as possible is documented, preferably automatically, in a system designed for this. These data must also be (or be made) findable, accessible, exchangeable and reusable[6]. Without this kind of foundation, artificial intelligence cannot build its model. The models make descriptive analytics and prescriptive analytics possible, thus elevating Total Talent Management from concept to reality.

How does it work?

Artificial intelligence can promote continuous improvement by linking a need to a concrete goal. With the goal as starting point, AI calculates optimisation scenarios. For this to work, the goals must not be formulated restrictively. Example: '*Maximise* the number of job applications for a hiring process' or 'Minimise the rate for a specific position without sacrificing quality'.

[5] https://www.reuters.com/article/us-amazon-com-jobs-automation-insight/amazon-scraps-secret-ai-recruiting-tool-that-showed-bias-against-women-idUSKCN1MK08G

[6] Fair principle: https://nl.wikipedia.org/wiki/FAIR_principes

A recruitment process goes roughly as follows.

Vacancy --> disseminate --> receive applications --> assess --> take on employee --> employment contract or temporary hire

If the *result* of every one of these steps is documented in the form of data - *input* - , it is possible to develop models. An example:

The *result* of *disseminating* an *HR vacancy* in the *Utrecht region* is *12 job applications*.

In this example, just two pieces of data provide for the result. One piece of data is fixed (HR vacancy). The second piece of data (Utrecht region) was decided by someone and is therefore changeable. The latter is referred to as an 'actionable variable'. Until now, one could manage this with a spreadsheet. Artificial intelligence can do more.

Practical example

Nétive VMS[7], market leader in the Netherlands for hiring systems according to 'Staffing Industry Analysts',[8] uses the 'Salesforce App Cloud[9]' Platform as a Service (PaaS). The Nétive system, a comprehensive 'Workforce Management System', has modules focused on both temporary hiring and the recruitment of permanent personnel:

Vacancy --> matching --> contracting --> compliance management --> performance accountability --> invoicing

'Salesforce Einstein' artificial intelligence can be linked to this system. Einstein automatically answers the following questions:

1. What happened? (descriptive analytics)
2. What has changed recently?

[7] https://www.netive.nl/

[8] https://www2.staffingindustry.com/site/Research/Research-Reports/Americas/VMS-Market-Developments-Summary

[9] https://developer.salesforce.com/platform

3. Why did that happen? (diagnostic analytics)
4. What could happen? (predictive analytics)
5. What is the difference?
6. How can I improve it? (prescriptive analytics)

This practical case involves the use of existing quantitative data on the hiring of freelancers and seconded workers who have two different forms of employment contract. The analysis was carried out twice: once with data up to 2017 and once with data from 2018 to Q1 2019. Both datasets were given the same goal: maximise the number of applications per vacancy.

Case 1: data up to 2017

Einstein proposed changing the number of hours per week worked by a contractor from 40 hours to 36 hours (prescriptive analytics). In that case the number of applications expected almost *tripled* with a stochastic variable of 1,400 vacancies. Einstein could designate other factors as 'influential', but not significantly.

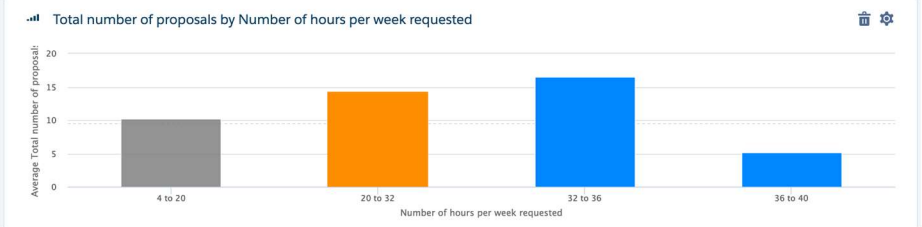

In 2017, the market in the Netherlands was not yet characterised by the current shortage. This outcome can also be explained by the fact that a freelancer devotes 4 hours per week to acquisition. Coming up with this explanation continues to be the work of humans. In this case, there are no research data available for this.

Case 2: 2018 – Q1 2019

The descriptive analytics show a different picture over time: the number of job applications for an ICT vacancy is declining significantly. Einstein concluded that the largest number of candidates is to be expected from secondment agencies, provided the option to renew an assignment is up for discussion:

The first bar in the diagram is the starting situation with the number of candidates from secondment agencies, the last bar is the predicted end result. The bars in between each represent a piece of data that influences the ultimate result. In this case, the changing labour market becomes visible; with the same goal as for case 1, the optimisation scenarios change over time. There is a growing scarcity of ICT talent.

Conclusions

Using artificial intelligence offers the possibility of a new form of working in this domain: Total Talent Management In this context, the influx of talent is documented in data which, combined with external market data and data available internally on talent, result in a single recruitment process model. This new working method creates the greatest increase in value for the organisation, namely insight into the value of human capital.

Points for attention

- Artificial intelligence requires ethical conduct, especially in the recruitment and hiring of workers.
- Artificial intelligence needs data and requires correct storage.
- Artificial intelligence is and can do more than search through large volumes of CVs faster, objectively and independently.
- It can make a difference if:
    - all relevant data and results of a recruitment or hiring process are recorded in applications specially designed for this.

- QCER is used to determine non-restrictive objectives (such as 'maximise the number of candidates').
- prescriptive analytics or descriptive analytics are used.
- Then artificial intelligence will be able to generate new insights and offer opportunities in the recruitment process, by constantly generating suggestions, taking into account changing factors of which no one was aware.

10. Self-driving cars

Victor de Pous

After rail transport via automated people movers, the self-driving car on the public road is no longer fiction but has become fact. Both in terms of ICT and law, we see concurrence. From a technical perspective, it mainly concerns Internet of Things applications, big data and artificial intelligence. For the legal aspects, it is the road traffic law that jumps to mind. The various legal frameworks for vehicle data are more in the background: privacy law, access to technical information under competition law - for maintenance, repairs and additional third-party services - and, for example, the installation of a black box in new cars envisioned by the EU. Zooming out, there is the threat of societal disruption if our automated fleet becomes virtually entirely dependent on technology from abroad, and if the supply falters.

Multifaceted legal framework

'The tax incentives for electric cars are an expensive measure to reduce CO2 emissions,' the General Audit Office wrote to the Lower House of Parliament on 26 June 2019.[1] Although members of the government contest the figures[2], there is no doubt that environmental law, tax law and subsidy law provide regulation for self-driving cars. On reflection, the robocar, regardless of energy source, illustrates the scope and diversity of legal and policy domains that can be applicable to digitalised physical objects.

Trials with a self-driving motor vehicle *on the public road* require amendment of the road traffic laws. That is happening. Since 1 July 2015, Dutch law allows the possibility of testing cars equipped with automated functions, such as adaptive cruise control, automatic parking and lane-keeping systems with the *driver*

[1] https://www.rekenkamer.nl/publicaties/kamerstukken/2019/06/26/fiscale-stimulering-van-elektrische-autos

[2] https://www.rekenkamer.nl/publicaties/brieven/2019/06/26/reactie-staatssecretaris-van-financien-en-staatssecretaris-van-infrastructuur-en-waterstaat-op-kamerbrief-fiscale-stimulering-van-elektrische-autos

in the vehicle.[3] The RDW grants an exemption for this. The Self-Driving Cars Experiments Act, which came into force on 1 July 2019, goes a step further and enables testing with a *driver at a distance outside the vehicle*, on the basis of a permit.[4] Aside from test situations, according to our current law, a car with automation level 2 is allowed to drive on the public road.[5] The vehicle can, in this context, take over certain driving tasks, such as adjusting speed and following the driving lane, while the driver's involvement in the vehicle remains necessary.

Liability and insurance

For liability, we can let the discussion begin as to whether a self-driving vehicle should perhaps have to take a driving skills test before driving on the public road. The answer seems to depend on who is considered the driver in the sense of the law (Article 61a RVV), but that is not necessarily the case. The RDW expects to grant the first 'S Driving Licence' to such a car at the end of 2019 or beginning of 2020, without, to our knowledge, the law undergoing any changes.[6] An updated practical test — and therefore driver training — could also possibly benefit road safety for the driver of such a car. That is also being worked on in the Netherlands.[7]

Article 185 of the Road Traffic Act provides extra legal protection to cyclists and other 'weak' traffic participants in the event of a collision with a 'strong' participant via strict liability for the owner or possessor of a motor vehicle *unless the latter can demonstrate circumstances beyond his/her control*. If a self-driving vehicle causes damage, the question arises of whether the owner or possessor can rely on circumstances beyond his/her control. If this is the case, we do not know,

[3] Decree of 15 June 2015 amending the Decree granting exemption for exceptional forms of transport (development of self-driving car)

[4] Act of 26 September 2018 amending the Road Traffic Act 1994 in connection with facilitating experiments with automated systems in motor vehicles.

[5] Autonomously driving cars are divided into levels, where level 1 stands for limited driving task assistance and level 5 for fully independent driving. https://www.automotivelectronics.com/sae-levels-cars/

[6] The Software Licence Driving Project (https://automotive-online.nl/management/laatste-nieuws/overig/22025-rdw-werkt-aan-rijbewijs-voor-zelfrijdende-auto-s). The question is whether this must take place every time, after certain software has been updated. And how does a new-style driving test deal with self-learning systems that develop over time?

[7] https://www.cbr.nl/nl/over-het-cbr/over/laatste-nieuws/nieuws/wie-heeft-de-controle-over-zelfrijdende-autos-2.htm

incidentally, whether the person can recover the damage from the producer or supplier of the AI system. Possible legal grounds for claims are attributable failure (contract law), unlawful act and/or based on the product liability regulations.[8]

From both a traffic law and insurance law point of view, the main rule is, for the time being, that the person behind the steering wheel of a car driving on autopilot remains the actual driver from a legal perspective, with the consequence that his/her hands must always be on the steering wheel.[9] Both the subdistrict court judge and Leeuwarden Appeal Court explicitly confirmed this principle in a criminal case against the driver who was holding a mobile phone while in a Tesla Model X driving on autopilot.[10] Speaking of insurance law: a Tesla owner also has to take out car insurance with the manufacturer (and have the vehicle serviced and repaired exclusively at a Tesla garage). This raises questions, both from the perspective of consumer law (limiting choice) and competition law (blocking off markets).

Other issues

In 1992, Technology Counsel at the Ministry of Justice T. de Graaf came up with something new: an innovative smartcard plan. Using a compulsory small computer installed in the car, which uses radiowaves to communicate with antennae in the road surface, vehicles could be monitored throughout the country and traffic offences, including speeding and parking offences and running red lights, could be registered. Automatic control would likewise take place for more 'administrative' offences, such as registration plate fraud, periodic vehicle inspection and road tax. Possibly to some extent because of fierce criticism from the *Stichting Waakzaamheid Persoonsregistratie* — the only privacy lobby group in the Netherlands at that time —, the idea remained a trial balloon.[11]

Twenty-five years later it emerges that many modern cars have a black box.

[8] The European Commission is working on a liability regulation for 'emerging digital technologies'. https://ec.europa.eu/digital-single-market/en/news/european-commission-staff-working-document-liability-emerging-digital-technologies

[9] https://www.consumentenbond.nl/autoverzekering/zelfrijdende-auto-verzekeren

[10] *SUSPECT VS PUBLIC PROSECUTION DEPARTMENT*, Utrecht District Court, 22 November 2018, ECLI:NL:RBMNE:2018:5707. In appeal: , Leeuwarden Appeal Court, 31 July 2019, ECLI:NL:GHARL:2019:6123.

[11] As cited in V.A. de Pous, *Privacy revisited; PrivacyWare 1992-1998*, Amsterdam, 2017.

The so-called event data recorder (EDR) is virtually de facto standard[12], the introduction of which apparently took place without much ado. The police, who can read out this data recorder for incidents and accidents and therefore obtain access to information on the driving situation from a number of seconds before, during and after the initial incident (usually the inflation of an air bag), are downright positive, as emerged earlier from a test in Rotterdam in 2016.[13] The European Commission and the European Parliament have since decided — from the standpoint of traffic safety[14] — that every car, van, bus and freight truck sold from 1 May 2022 onwards must have, among other things, an EDR system and 'intelligent speed assistance'. A regulation is being prepared.[15]

Vehicle data

The quantity of data that an automated car processes seems science fiction. All sorts of figures are circulating, in part because a number of different factors play a role, such as the number of sensors and cameras used and the degree of connectivity.[16] Even with a low level of autonomy, a 'connected car' generates and consumes some 25 Gigabytes of data per hour of driving. Speaking of big data. Google's car even produces almost *1 Gigabyte per second*.[17] The term 'vehicle data' is often used: a motley collection of data usually processed in real time, which can relate to the driver's driving behaviour, routes, location of the vehicle, traffic congestion, weather, road surface (slipperiness), technical condition of the vehicle, and more.

Legal aspects also accrue to these diverse data categories. If 'information about an identified or identifiable natural person' is concerned, European privacy

[12] According to the Dutch Association of Insurers, 'almost all cars have a sort of EDR'. https://automotive-online.nl/management/laatste-nieuws/schade/21900-verbond-van-verzekeraars-wil-verplichte-black-box-in-elke-auto

[13] http://www.ad.nl/binnenland/zwarte-doos-helpt-politie-na-autocrash~a4456e34/

[14]http://ec.europa.eu/transport/road_safety/specialist/knowledge/esave/esafety_measures_known_safety_effects/black_boxes_in_vehicle_data_recorders_en.htm

[15] https://eur-lex.europa.eu/legal-content/NL/TXT/HTML/?uri=CELEX:52018PC0286&from=EN

[16] Connectivity is one of the topics on which a discussion about standards arises. https://www.reuters.com/article/us-eu-autos-tech/eu-opens-road-to-5g-connected-cars-in-boost-to-bmw-qualcomm-idUSKCN1TZ11F

[17] https://www.kurzweilai.net/googles-self-driving-car-gathers-nearly-1-gbsec

legislation applies. The GDPR therefore applies to the data that the police read from an EDR system and also for much of the real-time information that the car constantly sends to the manufacturer, unless it is irreversibly anonymised or otherwise unable to be traced to individuals. In the event of technical information, competition law (market access) and consumer law (freedom of choice) are relevant, as is, for instance, software law, in particular the legal right to interoperability, on grounds of which the licensee should be able to connect other software to the software used (Article 45m(1) Copyright Act).

Digital threats

Another legal and policy aspect is the notorious dependency on ICT that makes us vulnerable — even society-wide. In the Netherlands Cybersecurity Picture 2019, the National Anti-Terrorism and Security Coordinator (NCTV) rightly calls attention to this in a general sense.[18] Because societal disruption is looming. What is striking is his broader view, which encompasses more threats than just espionage and crime. The Netherlands is (also) 'dependent on a limited number of providers and countries, which makes us vulnerable to their (changing) intentions', referring here to the US and China.[19] We stated earlier that virtually all our personal digital equipment - PC, laptop, tablet, smartphone - runs on American system software (OS), which is made available via a US email address.[20]

The question is: what is going to happen in the Internet of Things (IoT) segment, in which the automated car falls? That car manufacturers literally make their own digital vehicle operating systems is a *fait accompli*. Tesla presented its 'full self-driving computer' on 22 April 2019.[21] But it also concerns system software for all sorts of other and often linked functions, such as in-car entertainment systems, navigation, voice-operated software and more. Like Apple CarPlay and Android Auto, for instance. Is Europe losing the race here for this kind of software as well, which is moreover often supplied 'as a service'?

[18] https://nctv.nl/actueel/nieuws/2019/csbn-2019-ontwrichting-maatschappij-ligt-op-de-loer.aspx

[19] 'Most of the hardware and software are therefore designed or produced in China and the United States. Other countries can also have certain legislation that deviates from our privacy requirements or that results in access being obtained to the data of Dutch users or businesses,' the NCTV said.

[20] Three suppliers call the shots: Apple (Apple OS), Google (Android) and Microsoft (Windows). See *NEWSWARE*, year 32, 2018, number 6.

[21] https://electrek.co/2019/04/22/tesla-full-self-driving-computer-details/

Also striking is the NCTV's answer to the digital threats: 'resilience' is seen as 'the most important instrument to reduce risks because having any influence on the threat and dependency is too complex'.

Conclusion

Finding practical and generally accepted solutions for *civil-law* liability in relation to partially or fully self-driving cars may be difficult, but not impossible. A 'no-fault' traffic insurance comes to mind. The legal complexity increases, however, for issues concerning vehicle data, particularly with regard to ownership, access and use rights. Who has what rights? For the time being, the owner or possessor of the vehicle, given the practice, is in any event not *factually* in the best position.

Finally, our dependency on the car manufacturers and third parties who provide additional digital products and services is also increasing. On grounds of supplier dependency and other digital threats, the NCTV warns in a general sense even about societal disruption. The question is what resilience measures a Tesla driver can take against hacking or sabotage if the software is not updated on time with functional and security patches. What is he left with if the White House blocks the supply of US software to Europe in the context of a trade war?

Analyses

- Based on our law of contract, a precontractual duty to inform applies. Buyers of an automated car are entitled to information on, among other things, (i) the digital technology used (functionality, security, service level, update policy and more), (ii) the data that the car registers and needs and (iii) what happens with these vehicle data. (Privacy law also stipulates strict conditions for the processing of personal data, including the mandatory processing agreement.) *The question is whether sellers comply with this statutory duty to inform.*

- The European legislator advocates police being able to read out EDR data on grounds of traffic safety, to which argument the Dutch Association of Insurers adds the quicker settlement of claims and prevention of fraud. According to the sector, the police are a 'neutral party' in the event of an accident. *That is not necessarily the case, especially not if the EDR data expose criminal offences.* As a result, the driver could, as suspect, be

unwillingly cooperating with his own conviction, which, for example, is at odds with Article 6 of the European Convention on Human Rights.

- Broadening the focus here, the question of *ownership of* and, therefore, *control of vehicle data* is a common legal thread throughout the processing of vehicle data. Car manufacturers seem to overlook this by offering the owner or possessor of the motor vehicle few choices at present.

- Cheating car manufacturers have misled testing institutes, dealers, buyers and society on a large scale by the use of software codes with hidden functionality in the processing of CO2 emission data. *How reliable are the technical specifications of an automated car and how transparent is the manufacturer?*[22]

[22] A 2014 Tesla Model S 85 driver brought the manufacturer to court in the US in a class action on account of a software update which caused the maximum kilometre range of the battery to secretly decrease (RASMUSSEN VS. TESLA, INC.). Https://electrek.co/2019/08/08/tesla-owner-range-slashed-software-update-class-action-lawsuit/

11 Reflections on robo judges

Robert Sanders

In the seventh century before Christ, the Athenian legislator Draco codified the criminal law of his city state in order to curb class justice and arbitrariness. Tradition ascribes a proverbial strictness to Draco and his legislation: even minor offences, such as stealing an apple, incurred the death penalty, which, moreover, was carried out extremely consistently. This increased the predictability of the dispensation of justice: citizens knew in advance what to expect. But was this 'automation' of justice also justified and acceptable? People did not have computers or algorithms at the time, of course. Is today's fear of the draconian dispensation of justice as a result of the automation of court decisions justified? In and of themselves, because of their iterative nature and course of the proceedings, judicial proceedings lend themselves to a certain degree of automation: procedural steps and time periods are, after all, largely laid down in laws and regulations. Monitoring the time periods for the legal chess game of moves and counter moves could be automated without too much objection. That is not the case for the actual dispensation of justice, whereby the judge ultimately takes a decision.

Judge of the future

During his speech on the Day of Justice on 28 December 2017, F. Bakker, president of the Council for the Judiciary, anticipated the judge of the future and the influence of technological developments on the dispensation of justice. He sketched a picture of the judicial prototype of the top legal expert who renders top-quality legal decisions even in complicated matters, who interprets the law and safeguards the rule of law. For this kind of judge, in-depth *substantive* knowledge of algorithms and blockchain technology will become unavoidable. That is to say, he will need to be able to apply this knowledge in complex disputes about, for instance, smartphones or cybercrime. At the same time, Bakker says there will continue to be a great need for actual human judges who are there to hear and see the citizen. The judge of the human dimension, so to speak.

Central to Bakker's argument are, above all, the changing society and changing needs for the dispensation of justice. The danger lurks that the critical citizen will seek out alternative forums to resolve his dispute, causing the courts to lose ground. The courts will have competition and, when it comes to promoting customer satisfaction, there are no sacred cows as far as Bakker is concerned. If the predictability of justice can be increased by using software that analyses judgments (prediction algorithms), that should happen more. That means, therefore, that cases that lend themselves for this can largely be resolved using automation. Automated resolution which, as a joke, he refers to as the 'robo judge'.

Added value

Groningen professor of legal informatics and legal argumentation H. Prakken has doubts about this added value of artificial intelligence in the dispensation of justice.[1] According to him, the use of prediction algorithms does not bring automatic decisions in court cases any closer. After all, Prakken says there is an important difference between predicting and taking a legal decision: automating legal decisions for 'routine cases' would require a type of AI system very different from a prediction algorithm. Prakken turns it around: in routine cases, there is by definition an unequivocal outcome and there are no complex problems of evidence or interpretation. His robo judge is a human judge who is supported by technology. What is involved in that case is not a system that can predict the outcome of a case, but artificial intelligence that can apply legal knowledge and, in doing so, help substantiate decisions. According to Prakken, it is more realistic to aim for automated support for the human decision-maker using artificial intelligence so that humans and computers perform better together than a human or computer alone.

Legitimacy

This consideration of the added value of knowledge systems in the dispensation of justice gives rise to a more fundamental question, namely that of the *acceptance* of judicial judgments reached on the basis of algorithms.

The dispensation of justice is firstly an activity that, within the *trias politica* (separation of powers), constitutes an important component of our democratic rule of law. The term 'rule of law' clearly expresses that everyone in society, even the government, must adhere to the law. The law is the highest authority, therefore,

[1] H. Prakken, *Komt de robotrechter eraan?*', NJB 26 January 2018, pages 269-274.

and in legal disputes it is the job of the regular court to take decisions. The judge is part of the rule of law and, at the same time, is independent. The legal position of the judge is laid down in the law and he is not accountable to the other powers within the *trias politica*. The judge dispenses justice 'in the name of the King' and his judgments are, in principle, inviolable.

But what if this judgment is the result of an algorithm or other form of artificial intelligence? Such an algorithm must, of course, first be programmed. Is that done by a market or other party outside the justice system or by the government under the direction of a different power within the *trias politica*? Does the justice system in that case still have control of or influence on the software, even aside from the capacity to learn that characterises these systems? Higher forms of artificial intelligence, such as deep learning, can, after all, increasingly develop reasoning patterns that become entirely removed from our observation.

The fundamental question of the legitimacy of automated decision-making systems in the dispensation of justice has also attracted the attention of Minister Dekker (Legal Protection). In the budget deliberations in December 2018, the minister informed the Upper House of Parliament at length about how the government viewed the use of algorithms and artificial intelligence in the administration of justice.[2] In his letter to the Upper House, he formulated a number of interesting reference points against which the technological developments must be tested. Reference points that are implied in the constitutional and democratic values of our legal system.

According to the minister, of principal importance is that the balance within the *trias politica* is not jeopardised. The government (and by this he means mainly the executive power) has recently started using artificial intelligence in all sorts of administrative processes and the judicial system cannot lag behind, otherwise it will suffer in terms of effectiveness and legitimacy, that is to say, the role of the judge in the regular court will be granted less authority and the automatic acceptance of its judgments will diminish significantly.

Some principles of the rule of law can also come under pressure as a result of the use of artificial intelligence. This includes the right conferred by law (Article 17 Constitution) that no one can be kept away from the court and the five principles of the proper administration of justice from Article 6 of the European Convention on Human Rights and other international provisions. These are (1) *accessibility of*

[2] Letter from the Minister for Legal Protection dated 19 December 2018, Parliamentary Documents I 2018-2019, 34 775 VI, AH.

justice, (2) *internally and externally public,* (3) *trial within a reasonable time period,* (4) *independence and impartiality of the judge and* (5) *fair trial.*

Reference points

In order to retain these principles – he does not rule out a different form or interpretation – the minister outlined a number of reference points. The first reference point concerns the internal and external public nature of the administration of justice. There is no objection to the use of technical aids such as artificial intelligence and data analysis in order to bring together knowledge of the law: the judge is free, and sometimes obliged, to himself supplement the legal grounds. For the determination of the relevant facts, the judge is bound by rules and restrictions and he cannot, therefore, just pluck facts from the internet: not everything found there can be considered generally known facts. The 'robo judge' of the future will not base its judgment purely on Google search results. The facts must subsequently be determined and assigned value. The judge will have to render a moral judgment based on and about those facts.

A second reference point identified by Minister Dekker is therefore related to the inevitable interpretation of law and facts: there is often not just one correct solution to a legal issue or case. The legal debate has independent value in the courtroom and the question must therefore be asked of to what extent the use of artificial intelligence leaves room for that legal debate.

How an AI system then reaches its result is outside the bounds of human observation. For the rest, a parallel can be drawn with the secrecy of the judge's chambers in this respect. For this reason, judges are required to provide their judgments with sound, i.e. adequate and convincing, substantiation. According to Minister Dekker, this substantiation cannot therefore be that the decision is simply the inevitable outcome of a statistical analysis of the case law up to that point.

Finally, the minister presents the reference point of the independent value of settlement by human instead of automated judges. As Frits Bakker also noted, someone seeking justice wants to be heard. So-called procedural justice, the acceptance of judges' decisions, is based not only on the content of the decision, but just as much on the treatment of the citizen seeking justice.

Conclusion

Opinions diverge on the question of whether artificial intelligence can take over the role of the human judge and whether a knowledge system based on algorithms can function as a 'robo judge'. It seems inevitable that in the future, higher and higher

demands in terms of knowledge of technology will be put on the judge's role as we know it now. And there is little doubt that technology in the form of knowledge systems can assist the judge in that respect. The examples of such knowledge systems are already available and if it were up to Frits Bakker, there would be 'no sacred cows' when it comes to reserving terrain for the judge. Reserving terrain in a changing world in which the regular court is not the only forum being looked to for dispute resolution. And yet, at the moment these changes are not yet to the point that a judicial decision spat out by a computer will be automatically accepted. Procedural justice, which seeks to give the party seeking justice the sense that he and his interest have been heard and seen, but also the importance of the legal debate and transparency on how a decision has been reached stand in the way of a 'robo judge'. That is to say, in the form of a 'robo judge' that takes over all facets of the role of the actual human judge in a regular court.

Analyses

From the foregoing bird's-eye view reflections, we can, against the backdrop of this discussion, identify a number of basic assumptions:

- When talking about 'robo judges' or 'robo justice', a distinction must be made between rule-based knowledge systems, process descriptions and case analyses on the one hand, and intelligent moral decision-making systems on the other. Only the latter form of applied artificial intelligence has the potential to come close to replacing the judge as we now know it.

- As an open door: the development of technology and artificial intelligence will inevitably impact the administration of justice, even if just in the sense of the legal knowledge and practical expertise that will be demanded of judges in relation to technology and robotisation. The justice system must guard against becoming the weakest link in this specific area in the *trias politica*. Within this *trias politica*, the legislator also still has steps to take, incidentally. I am referring in particular to individual members of Parliament who are increasingly faced with complex and technical legislation, such as the General Data Protection Regulation, and are nonetheless deemed to be able to assess legislative proposals from that perspective as well.

- Besides what is technically possible, the introduction of artificial intelligence also constitutes a threat to the legitimacy of the administration of justice. Although the use of technology - and artificial intelligence as well - in the administration of justice is supported and encouraged, the principles of our rule of law set certain limits to this application.

- A part of this legitimacy and the acceptance of the administration of justice is 'procedural justice'; in other words: the human dimension in which the party seeking justice is given space to feel heard. No single legal case is entirely identical to another. And every case is open to debate. That, too, constitutes part of our sense of justice. There must always be room for reasonable doubt or the circumstances of the case. There is still a long way to go before proceedings before an actual 'robo judge' can satisfy these requirements. Let us not forget that the Draconian law was largely rolled back within a few generations.

12. Towards legal personality for robots?

Natascha van Duuren

Large parts of European law, and Dutch law, stem directly from the Corpus Iuris Civilis: the law of the Romans. In those bygone times, there were no robots, of course, let alone robots that functioned autonomously. Roman law pertained primarily to people of flesh and blood, therefore, as bearers of rights and obligations. As a person having legal rights and with legal personality.[1] The term 'legal entity' or 'legal person' did not appear until the 19th century when, in legal transactions, one also had to deal with 'communities' of people. Besides natural persons, the law also recognises a legal person as a person having legal rights. Animals are not persons having legal rights, even though some argue for animals to be granted rights.[2] A human can, however, be the legal owner of an animal and can be held liable for its conduct. Robots are now rapidly integrating in our society. This raises the question of whether robots should be given a separate status in our legal system. Saudi Arabia, in any event, thought they should, and granted humanoid Sophia citizen rights in 2017. Virtually certainly a first worldwide.

What is a robot?

In 2012, the European Union came to the realisation that it is inevitable that robots will have a place in our society. The 'RoboLaw' project was started,[3] with the goal of identifying the legal and ethical implications of the rise in robotics technologies and examining whether the existing legal framework is adequate and feasible. Before the results are discussed, we will first devote some attention to the object of study. The literature contains countless definitions, ranging from descriptions pertaining to industrial robots to humanoids, which, like Sophia, resemble a

[1] The fact that natural persons are bearers of rights and obligations has since been incorporated in Article 2 of our Constitution and Article 1:1 of our Civil Code.

[2] www.dierenrecht.nl

[3] www.robolaw.eu

human.[4] RoboLaw uses the following definition: '*an autonomous machine that displays physical similarities with humans, consists of sensors, power supply and operating elements, and is capable of autonomously performing certain tasks in interaction with the environment or humans.*'

RoboLaw Project

In response to the RoboLaw project, the *Guidelines on Regulating Robotics* were drafted. These guidelines advise, among other things, that Europe-wide rules be established for liability and insurance law. There was no mention yet of separate legal status for robots in our legal system.

Several years later, the European Parliament adopted a resolution.[5] According to it, the current legal framework does not suffice for legal liability for damage caused by a robot in scenarios in which robots can take autonomous decisions. It therefore argues for new, contemporary rules. Rules in which a new legal category is created for robots.

Open letter

Not everyone agrees with this recommendation. In 2018, a clear opposing opinion was given by a number of political leaders, European Artificial Intelligence and Robotics Experts, captains of industry, medical and ethical experts in an open letter to the European Commission.[6] They stated that they advocated a civil-law framework for the liability of robots operating autonomously. However, separate legal status for robots would, according to them, result in ideological and pragmatic objections and would furthermore be nonsensical. The European Economic and Social Committee also voiced a critical opinion in 2018.[7]

Natural person, legal person or separate qualification?

Back to the question: is legal personality for robots desirable? Opinions differ strongly on the answer to this question and there has been much public debate.

[4] The robot Sophia was given citizenship of Saudi Arabia on 25 October 2017. As such, Sophia has the same rights and obligations as a man.

[5] Resolution of the European Parliament of 16 February 2017 with recommendations to the Commission on civil-law rules in relation to robotics (2015/2013(INL))

[6] Open letter to the European Commission on Artificial Intelligence and Robotics, www.robotics-openletter.eu

[7] www.eesc.europa.eu/en/our-work/opinions-information-reports/opinions/artificial-intelligence.

For instance, also in the disagreement printed in the NRC between university lecturers from two Dutch universities.[8] While one argued that there is nothing wrong with developing and revising a legal framework for AI and robotics and that this evidences a forward-looking view, the other argued that the question is what problem we, as society, actually want to solve with this. According to him, it would be good to think about issues concerning liability. However, it would not be a good thing and would be undesirable to sketch out a future in which robots could start legal proceedings against humans.

In order to take a position in this discussion, it is important to determine what it concretely means for robots to have legal personality.

Legal personality for robots means that they can independently take part in legal transactions. For example, they can conclude contracts, but also cause damage and be held liable. Legal personality could, of course, only be an issue for robots that operate autonomously without human instructions or supervision and not for robots that follow the instructions of humans or businesses.

Based on our current law, natural persons and legal entities have legal personality. The fact that a robot cannot be equated with a natural person should be immediately clear to everyone. At the same time, in the literature, the (in my view not irrelevant) point is broached that Dutch law does not contain a definition of a natural person. The law only states 'all who are in the Netherlands are free and entitled to enjoy civil rights'.[9] The relevance of the question of whether a robot can or cannot be equated with a natural person arises when we realise that there are many natural persons who have (one or more) artificial body parts. Furthermore, not all natural persons are 'entitled to enjoy civil rights'. Take minors, for instance, or people placed under guardianship. There can also be arguments against criteria such as 'free'. After all, animals also have free will.

In Saudi Arabia, it was decided that the robot could indeed be equated with a natural person. In 2017, a robot named Sophia[10] was granted citizenship. Sophia has the same rights as a Saudi man. The robot is a social humanoid robot with a female appearance modelled after the actress Audrey Hepburn. Sophia is capable of recognising human faces and recognising emotions, but also has her own emotions. Japan, too, has granted citizenship to a robot: chatbot Shibuya Mirai[11].

[8] www.nrc.nl/nieuws/2018/04/20/moeten-robots-een-aparte-juridische-status-krijgen-een-twistgesprek

[9] Article 1:1 DCC.

[10] www.hansonrobotics.com/sophia

[11] www.shibuyabunka.com

Legal entities

Another question is whether a robot can be equated in our legal system with (non-natural) legal entities. In the open letter to the European Commission mentioned above, objections were presented against both the qualification of a robot as a natural person and as a (non-natural) legal entity. According to the writers of the open letter, regarding a robot as a natural person comes up against the undesirable consequence that a robot could then also be granted (fundamental) rights. The right to human dignity and human integrity are mentioned as examples.

The qualification as legal entity comes up against the fact that this would presume that there is always a natural person behind the robot, according to the writers of the letter. (Just as there are always one or more directors 'behind the legal entity', who can be called to account.) This is not necessarily always the case for a robot operating autonomously. In my view, this is a genuine objection.

Another option would be to include a separate qualification for robots in the law. In that case, a third type of person having legal rights would actually be added to the law. This is not such a crazy idea. After all, robots will become more and more integrated in our society and in our social interactions. Furthermore, in the case of autonomously functioning robots, it will not always be so simple, if even possible, to determine who can be held responsible for the actions (or omissions) of the robot. The current legal system would be in a jam in that case. After all, an autonomously functioning robot cannot bear legal responsibility if it does not have legal personality. An argument against legal personality is, in my view, that this could have extremely undesirable consequences, that the developers of the robot could easily shift their own responsibilities and liabilities onto the robot that they have created.

Europe seems now rightly to want to make haste with providing for a suitable ethical and legal framework for AI.[12] In mid-2019 at the latest, it will publish a report on the possible lacunae in current security and liability frameworks. The question is whether, in doing so, it will address the desirability or undesirability of legal personality for robots.

In conclusion

At the moment, the question of whether legal personality for robots is desirable is only relevant for a very limited number of robots. Specifically, this concerns robots

[12] www.eumonitor.nl

that operate autonomously and that perform social and economic functions. It will not be long before this question applies for a larger group of robots. The haste Europe seems to want to make at present is therefore justified.

Considerations

- In the near future, numerous robots will be integrated in our society, whereby it will not be so easy to determine what person or company is 'behind the robot'. That is an undesirable situation, of course. Clarity must be provided on this, even if only from the viewpoint of transparency.

- The question is, moreover, who can bear legal responsibility for the actions (or omissions) of this robot. This question will mainly be relevant for robots that operate autonomously and that perform social and economic functions.

- The answer to this question is complex, beyond that it is not merely a legal question, but rather a moral and ethical question.

- However, is granting robots legal personality the only way to answer the question of legal responsibility for the actions or omissions of the robot?

- Given the rapid developments in the area of robotics, the European Union has rightly indicated it wants to make haste in identifying possible lacunae in current security and liability frameworks. Whether, given the broad complexity of this question, it will also make a pronouncement on the desirability of legal personality for robots is very doubtful.

13 Governance

Joan Baaijens, Leon Dohmen and Liesbeth Ruoff-van Welzen

The question is not whether we want our society to be digitalised, but how to do it decently. AI Governance can play an important role in this, but this means for structure and integrity is still largely unknown territory. There is not even a broadly supported unequivocal definition or standards framework. Nor does the term currently appear in the online encyclopaedia Wikipedia, or in the Gartner IT Glossary. Particularly striking in relation to the topic of artificial intelligence are the concerns about abuse and the infringement of human rights, for instance, about abuse of power, privacy violations, loss of employment and injustice. These concerns indicate without a doubt that attention to and further development of AI Governance is necessary.

Distinction between IT and AI

According to the Gartner IT glossary,[1] IT is 'the common term for the entire spectrum of technologies for information processing, including software, hardware, communications technologies and related services'. Over the past decades, IT has developed from an instrument for improving efficiency into a multifaceted instrument with dynamic value perception.[2] According to Gartner, AI 'applies advanced analysis and logic-based techniques, including machine learning, to interpret events, support and automate decisions, and take actions'. AI is inextricably connected with artificial neural networks, big data, algorithms and autonomous systems. With IT, the person determines the result of a software program. With AI, the system can become a so-called 'black box' that decides and acts independently based on an inherent technical and mathematical logic. Loss of control, insufficient transparency and specific legal problems such as liability are new risks. That is why AI is fundamentally different from other manifestations of IT.

[1] https://www.gartner.com/it-glossary/

[2] https://hal.inria.fr/hal-01448051/

Current state of AI

The term digitalisation refers to a collection of technologies, such as robotics, artificial intelligence, algorithms and big data. This last term generally refers to: a combination of datasets consisting of large volumes, with a high degree of flexibility and variety, that enables new forms of processes and analyses for users and organisations. Together they are creating a new wave of digitalisation of which AI is part.

The physical world is increasingly being represented virtually. This creates an intrinsic intertwining between this physical and virtual world and constant feedback between the physical and virtual realities. Products and/or services are developed and adapted based on the analysis of digital data (see figure[3]).

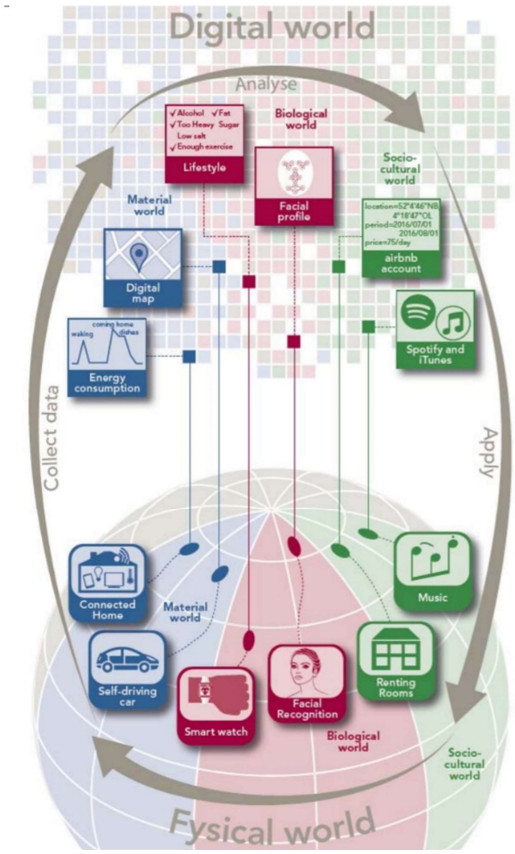

[3] https://www.rathenau.nl/sites/default/files/2018-02/Opwaarderen_FINAL.pdf

On the one hand, the great potential of AI is pointed to. Intel, for example, provides an inexpensive artificial neural network to stimulate AI developments. On the other hand, the dangers of AI are mentioned as well. Ranging from killer robots to a lack of transparency about how an autonomous system reaches a decision and the risks of bias that could occur in this context. [4]

A smart (and tempting) combination of cloud and other services and the unbridled collection of personal data have resulted in, among other things, a limited number of large data monopolists. In their wake, other organisations and businesses have also discovered the value of collecting personal data with just one goal: the commercial exploitation of these data. The data monopolists and data dealers do not take the legislation focused on protecting personal privacy terribly seriously in this context. Data are processed into profiles and sold to telecom providers and banks.[5] It can happen that a person may be unable to obtain a mortgage on the basis of this profile. 'People in financial difficulty could easily be tempted to sell or give away their data in exchange for the use of services. In an extreme form, it could result in what is referred to as "data prostitution".'[6] In the continuing digitalisation and development of AI, earning (a great deal of) money appears to prevail over respecting human rights.

From IT Governance to AI Governance

In the beginning (1993), IT Governance was a means of providing structure to the set-up and use of IT and attuning this to the company's objects.

Because of the introduction of new technology such as the internet, email and social media, the focus was increasingly pointed outwards and the emphasis shifted to the decent use of IT. As a derivative of Corporate Governance, IT Governance increasingly developed into a societal issue.[7] We must remember here that every form of governance refers to an underlying institutional structure in which the integrity of a transaction is determined.[8] Governance therefore concerns the reliability and management of transactions, such as the configuration of

[4] https://www.blogit.nl/vloek-zegen-kunstmatig-neuraal-netwerk-79-us-dollar/

[5] https://nos.nl/artikel/2288160-kabinet-burger-moet-wraakpornobeelden-makkelijker-offline-kunnen-krijgen.html

[6] https://kennisopenbaarbestuur.nl/media/255845/verslag-conferentie-mydata-2018.pdf

[7] https://www.blogit.nl/governance-intern-organisatiethema/#_edn4

[8] The Mechanism of Governance, O.L. Williamson, 1996

information facilities. In this context, 'Internet Governance' can be regarded as regulating internet use.

Corporate and private actors

For all forms of governance issues, it is very important to make an analytical distinction between basic forms of operating actors, specifically in terms of a 'corporate' actor and a 'private' actor or natural person.[9] A 'corporate' actor is a form of social organisation in which a natural person (private actor) has transferred his or her right to act to a formal legal person on the basis of a contract. Corporate actors usually act on the basis of their (defined) interests. The interests of natural persons are not safeguarded in advance in this context. For instance, the interests of natural persons, in terms of private life (privacy, confidentiality, transparency and retaining one's identity) can become damaged through information transactions because of the mode of conduct of the corporate actor. The distinction between societal interest and private interest can therefore be related to the difference in conduct between a corporate actor and a natural person (see the gas extraction in Groningen and the damage suffered by individual citizens). AI Governance also has to do with this fundamental distinction. And therefore with the serving of the societal interest in general and respecting individual human rights in particular.

No standards framework

For IT security issues, the ISO27001 standard is a generally accepted standards framework. There is no such standards framework for AI as yet. On 7 February 2019, the NEN organised a kick-off meeting to arrive at a standard for building AI applications. The NEN estimates that this standard will not be available for another several years. One consequence of this is that businesses are developing quality marks themselves. One of these was central to the discussion at the round-table session organised by the KNVI on 27 March 2019. Over the past several years, codes of conduct have been developed by various parties. Some of these focus specifically on AI. A distinction is made here between building the algorithm itself (the 'engineering' work) and what an organisation does to guarantee that responsible and transparent algorithms are being built. The latter explicitly refers to governance. A distinction is made here between technology, ethics and social/legal frameworks.[10] In the absence of a framework of standards and (adequate) AI

[9] See: Power and the Structure of Society, James S. Coleman, 1974.

[10] https://dash.harvard.edu/bitstream/handle/1/34390353/w6gov-18-LATEX.pdf?sequence=1

governance, in the technical and ethical context, the competencies and moral awareness of the builders of AI solutions play an important role in the identification of (potential) human rights violations and the protection of these rights.

Visible and invisible AI Governance

Governance consists of intentional (visible) agreements and rules. There is also spontaneous, invisible, governance. Invisible governance is driven by standards and values, such as tolerance, and the informal organisation. Visible and invisible governance consist of agreed respect.[11] For the time being, AI Governance is primarily visible in codes of conduct and legislation such as the GDPR. AI governance is, as such, mainly visible via a legal framework, such as the law of contract. Guaranteeing the building of responsible algorithms seems then mainly to involve encouraging algorithm builders to put the economic interest of their own organisation (the risk of a high fine) above the societal interest of protecting individual human rights.

The ethical framework, represented by codes of conduct, loses out in this process. Within the technical framework of AI Governance, builders of AI solutions often work together as part of informal organisational forms. Because organisations neglect to do so, builders must hold each other accountable to a much greater extent via the alternative of invisible governance, which prioritises tolerance and respect for human rights. After all, their own human rights are also involved. Builders must be more aware of the impact of AI solutions. This can be achieved by asking questions such as:[12]

- Who will lose and who will gain from this solution?
- Who will probably experience the most damage from this solution?
- What unintended side effects could this solution have?

To be able to better guarantee the above, the professionalisation of builders of AI solutions is increasingly needed.

---

[11] http://petersvmd.com/PrincipalAgent/short%20principal%20agent/Visible%20Invisible%20Governance.pdf

[12] https://www.rathenau.nl/nl/digitale-samenleving/knvi-geef-iters-de-ruimte-om-schending-grondrechten-te-benoemen

Competencies

It is a legitimate question to ask what level of knowledge and skills can be found on the board of directors and supervisory board, since these are ultimately responsible for the governance in the organisation, in this case concerning AI. Do they have sufficient knowledge or have they invested enough in themselves to shape proper AI Governance within the organisation, whereby there is (more) balance between economic interest and respect for human rights?

The National Register and Nederland ICT recently presented a 'baseline measurement'[13] focused on the question: 'How do managements/boards of directors, supervisory boards and other supervisors of listed funds, mid-market companies and semi-public organisations view or deal with current issues relating to digitalisation?' Knowledge components[14] that emerge in this case include: knowledge about new technologies and techniques for determining the risks and opportunities presented by technology, but also being able to identify the key milestones within a proposed plan. Skills that emerge are in areas such as:

- Review and analysis of the effects of implementations of new technology;
- Understanding of the impact of new technologies on the business;
- Understanding of the benefits of new technologies for the business and how these can add value and yield a competitive advantage, but also how they can pose a risk;
- Understanding of the legal and regulatory landscape.

In and of itself, this knowledge and these skills should no longer be lacking on a management board or board of directors in the 21st century. And yet most in this group say they do not have it. The consequence, therefore, is the lack of adequate governance for the development of AI solutions.

Conclusion

As a sub-area of the continuing digitalisation, a number of technologies together make it possible to build AI solutions. AI offers much economic and societal potential. But at the same time, the risks of autonomously operating systems that are artificially intelligent are pointed out. Adequate AI Governance not only takes

[13] Baseline measurement on digital transformation in board rooms in the Netherlands by Prof. Marcel Thaens and Prof. Valerie Frissen, carried out in 2018 and published in 2019.

[14] http://www.ecompetences.eu/wp-content/uploads/2014/02/European-e-Competence-Framework-3.0_CEN_CWA_16234-1_2014.pdf

into account the organisation's own profit interest, but also respects individual human rights. With visible governance, the legal framework is dominant. With invisible governance, matters such as trust and acquired standards and values play a role. This can be counterbalanced by the builders of AI who can hold each other accountable within informal organisational forms (invisible governance) for the advantages and disadvantages of the AI solutions they build.

Considerations

- Businesses and the government have a duty of care to protect people's fundamental rights in the further digitalisation and development of AI. In the continuing development of AI, there is a danger of lack of respect for these human rights (for example privacy) and abuse of power by organisations (corporate actors). Is this a lack of moral sense on the part of organisations?

- Because organisations do not have their AI Governance in order or sufficiently in order, builders of AI solutions (private actors) must be more aware of the impact of their solutions. Can a code of conduct provide support in this respect?

- Builders can themselves, regularly and at various moments in concrete work situations, broach critical questions about the AI solutions they build (invisible governance).[15] Competencies, moral awareness and professionalisation are needed for this.

[15] https://www.rathenau.nl/nl/digitale-samenleving/knvi-geef-iters-de-ruimte-om-schending-grondrechten-te-benoemen

14 Privacy-law aspects

Jeroen van Helden

We learn by making mistakes, is how the saying goes. This applies for a self-learning AI application as well. The logic of a self-learning AI application is not pre-programmed, but is given shape within an artificial neural network that is fed with examples, from which the system then learns. These examples can contain privacy-sensitive information. At a certain point, the application has learned enough; in other words, it is 'mature' enough to go 'live' and can actually start playing a role in the real world. The application may be used to operate a car, for instance, in which new data are in turn collected. Certain parties involved in the development of self-driving transport, such as Alphabet with subsidiary Waymo (Google), are suspected of ultimately being less interested in selling self-driving cars than in collecting the data that are generated about the driver and cashing in on those data. An AI application can also be used to take decisions that have a certain impact on a person's private life, for the selection of potential defrauders, for instance, or in screening job applicants. So it is high time to examine the privacy-law aspects of self-learning applications.

Framework

Any discussion of privacy will quickly arrive at a mention of the General Data Protection Regulation (GDPR). The GDPR took effect throughout the European Union on 25 May 2018 and provides rules for the (automated) processing of personal data. This type of privacy regulation arose in the 1970s when it became clear how easy it was to store and further process large volumes of data using computers. The GDPR is meanwhile quickly becoming the international standard for regulations on data processing.[1]

Because of the attention garnered by the GDPR, one might think that this European law is the only privacy-law framework which society needs to deal with in the development and application of artificial intelligence. That is not entirely

[1] For instance, the California Consumer Privacy Act (CCPA) which will come into effect on 1 January 2020 is partly based on the GDPR.

correct. Privacy as a fundamental right has long been anchored in numerous international treaties, as well as in the Dutch Constitution. The fundamental right to the protection of our private life encompasses much more than the rules on processing personal data and can emphatically also be relevant for AI applications. This article first discusses AI in relation to the GDPR, before discussing AI in relation to the fundamental right to privacy.

Datasets

Whether an AI application is trained on the basis of supervised learning, unsupervised learning, reinforcement learning or some other learning method, data are always needed to train and develop the automated system, so that it can subsequently go to work independently. The GDPR interprets the term 'personal data' broadly, so that the GDPR will apply to many datasets.[2] This presents the following challenges:

- Based on the GDPR, personal data can only be collected for well-defined, explicitly described and justified purposes and the data may not be subsequently further processed in a manner inconsistent with those purposes. Data subjects must also be informed about that use and their permission may be required. The data system behind an AI application is often complex and seldom transparent. Aggregate and enriched data originating from multiple sources are used. In practice, the obligation of purpose limitation and duty to inform are therefore perceived to be difficult or even impossible to fulfil.[3]
- According to the GDPR, personal data must be sufficient, fit for purpose and necessary for the purposes for which they are processed. The development of AI applications often benefits from large datasets. The trick is maximal data processing rather than minimal data processing. Although some techniques, like generating synthetic data or using Federated Machine Learning, can help limit the processing of personal data, there is no simple solution that

[2] With regard to personal data and anonymisation, see Article 29-WG, Opinion 4/2007 on the concept of personal data, 20 June 2007, and Article 29-WG, Opinion 5/2014 on Anonymisation Techniques, 10 April 2014.

[3] For example, see: Licht op de digitale schaduw: verantwoord innoveren met big data [Light on the digital shadow: innovating responsibly with big data], Report from the expert group on Big data and privacy to the Minister of Economic Affairs, August 2016, page 21.

eliminates the fundamental tension between the principles of the GDPR (few data) and AI applications (many data).[4]

When using AI applications, one must also be aware that AI applications are capable of discerning patterns in the dataset that humans would not be able to discover. As a result, 'innocent' data can become 'sensitive' data in the hands of an AI application. For instance, an AI application could be able to conclude a person's political leaning or medical condition from what appear to be innocent data. It is also conceivable that an AI application could convert what appears to be anonymous information into data identifying an individual by name. Recent research also shows that self-learning AI applications do not quickly forget rare and sensitive training data and that they can unintentionally produce these data in live environments.[5]

Automated decision-making

Automated decisions are being taken all around us. For example, the question of which advertisement you are shown online or the amount of the supplement to which you are entitled according to the tax authorities' website. The advantages are unmistakable: speed and consistency. It is suspected that most of the decisions in the sense of the General Administrative Law Act (Awb) are taken via an automated process.[6]

The GDPR contains a few specific rules on automated decision-making. First of all, an individual must be able to know that automated decision-making is being used. That is why the controller has the obligation to inform the data subject about this.[7] Part of this duty to inform is that the data subject must be informed about the 'underlying logic' of the application. This does not require that the algorithm's

[4] The use of these and other technologies are recommended by the Council of Europe, *Guidelines on Artificial Intelligence and Data Protection*, 25 January 2019, and the International Working Group on Data Protection, *Working Paper on Privacy and Artificial Intelligence*, 64th meeting, 29-30 November 2018.

[5] Carlini, N. (2019). Evaluating and Testing Unintended Memorization in Neural Networks, https://bair.berkeley.edu/blog/2019/08/13/memorization/.

[6] van Eck, M. (2018). Geautomatiseerde ketenbesluiten & rechtsbescherming: Een onderzoek naar de praktijk van geautomatiseerde ketenbesluiten over een financieel belang in relatie tot rechtsbescherming [Automated chain decisions & legal protection: A study into the practice of automated chain decisions on a financial interest in relation to legal protection], page 30.

[7] Art 13(2)(f) GDPR and Art 14(2)(g) GDPR

functioning be explained in (technical) detail or that the algorithm be published. It must be made clear in a comprehensible manner, however, how the AI application works and on the basis of what criteria a decision is arrived at.[8]

The data subject also has the right not to be subjected to a decision based exclusively on automated processing which has legal effects for him or which otherwise affects him substantially.[9] Although this article is formulated as a 'right', according to the European privacy regulators, there is a de facto general prohibition on fully automated decision-making, notwithstanding exceptions.[10] The termination of an employment contract exclusively on the basis of an automated decision is therefore prohibited. Targeted advertising based on profiling will generally not have a significant impact on the data subject and consequently does not generally fall under a prohibited form of automated individual decision-making.[11]

Automated decisions are usually reached on the basis of pre-programmed decision rules. In that case, it is relatively easy to inform users about the 'underlying logic' of the system and to monitor whether the system has arrived at a qualitatively good decision. It can be more difficult to satisfy these obligations if an AI application, trained on the basis of data, has taken the decision.

---

[8] Article 29-WG, Guidelines on Automated individual decision-making and Profiling for the purposes of Regulation 2016/679, 6 February 2018, page 25.

[9] This right is contained in Article 22 GDPR. Incidentally, the right does not apply to an unlimited extent. The right does not exist if i) the decision is necessary for the establishment or performance of an agreement between the data subject and a controller; ii) the decision is permitted pursuant to a statutory provision that applies to the controller and that also provides for appropriate measures protecting the rights and freedoms and justified interests of the data subject; iii) the decision is based on the explicit permission of the data subject; or iv) the decision, other than on the basis of profiling, is necessary in order to satisfy a statutory obligation borne by the controller or one that is necessary for the fulfilment of a task carried out in the public interest. In these cases, several additional conditions may apply, such as that the data subject has the right to make his/her position known and that the automated decision cannot be based on any special categories of personal data.

[10] Article 29-WG, Guidelines on Automated individual decision-making and Profiling for the purposes of Regulation 2016/679, 6 February 2018, page 19.

[11] Ditto, page 22. The situation could be different under certain circumstances, however. For example, if prices are adjusted in a targeted manner, making a particular product or service unaffordable for an individual.

In that case, the relationship between the input and the ultimate output (the decision) is more difficult to trace – including for the developer of the application. After all, the decision-making becomes a kind of black box.[12] For example, an AI system developed by an insurer decided that drivers of red cars should have to pay a higher premium than drivers of cars of other colours.[13] This brings me to the second legal framework.

Privacy as fundamental right

The European Court of Human Rights (ECtHR) interprets the right to privacy as contained in Article 8 of the European Convention for the Protection of Human Rights and Fundamental Freedoms (ECHR) broadly. Based on that, citizens not only have a certain right to be left alone by the government, but they also derive from that a right to personal development. The ECtHR also tends to draw other fundamental rights within the scope of Article 8 ECHR, such as the right not to be discriminated against and the freedom of speech.[14] Finally, the right to privacy is not only relevant in constitutional and administrative law (vis-à-vis the government), but the right can also inform private-law legal relationships (between citizens and businesses), the so-called horizontal effect of fundamental rights.

We have now discovered that many datasets are not neutral, but reflect societal prejudices and socio-economic inequalities. An AI application quickly adopts this predisposition, possibly in reinforced form. Amazon, for instance, used an application to assess CVs but it emerged that the self-learning program had a conspicuous preference for men.[15] This problem is known as 'garbage in, garbage out' or 'bias in, bias out'.

The European privacy regulators expect controllers that use AI applications to frequently audit their algorithms and datasets for possible prejudices and other discriminatory elements that could be contained in them.

Finally, AI technologies, like many other digital technologies, often use insights derived from neuroscience and psychology with the aim of getting users

[12] Castelvecchi, D. (2016). Can we open the black box of AI? *Nature*, 538, pages 20–23.

[13] Interview with E. Haasdijk, National AI course, track 6.

[14] van der Sloot, B. (2015). Privacy as Personality Right: Why the ECtHR's Focus on Ulterior Interests Might Prove Indispensable in the Age of "Big Data". *Utrecht Journal of International and European Law*, 31(80), pages 25–50.

[15] https://www.reuters.com/article/us-amazon-com-jobs-automation-insight/amazon-scraps-secret-ai-recruiting-tool-that-showed-bias-against-women-idUSKCN1MK08G

to spend time and money in a virtual environment. Examples are the automatic modification of content based on profiles and the Fear Of Missing Out (FOMO) principle. It is not inconceivable that the use of such technologies could, under certain circumstances, be in violation of the right to privacy as interpreted by the ECtHR. The alternative is also conceivable: that, on the basis of Article 8 ECHR, a person could claim a right to access to a friendly and helpful care robot.[16]

Conclusion

Self-learning AI applications thrive best in data-rich environments. It is therefore no surprise that this category of artificial intelligence raises countless privacy-law questions, especially since the regulations on data protection were themselves introduced relatively recently. To some extent, the problems are the same as the problems of any big data collection. Where AI applications are to take decisions themselves, additional questions arise, for instance about how the duty to inform is fulfilled and how the rules for automated individual decision-making are applied. Depending on the application and the circumstances of the case, issues relating to fundamental rights (privacy) can also arise.

Points for attention

- Self-learning AI applications process large volumes of personal and other data so the GDPR must remain a constant point for attention in the development of these applications. There are no simple solutions; in a general sense, the advice is therefore to always make a careful, expert and documented consideration.

- The GDPR stipulates specific obligations for the use of automated individual decision-making. It is relatively more difficult to satisfy these obligations with self-learning AI applications than with the use of traditional applications based on decision rules; however, it is not impossible.

[16] This topic is explored in Van Est, R. & Gerritsen, J.B.A. (2017). Human rights in the robot age: Challenges arising from the use of robotics, artificial intelligence, and virtual and augmented reality – Expert report written for the Committee on Culture, Science, Education and Media of the Parliamentary Assembly of the Council of Europe (PACE), The Hague: Rathenau Institute, pages 18-26.

- The fundamental right to protection of one's private life is interpreted broadly by the European Court; consequently, this right can also be relevant for the application of artificial intelligence. For example, if an AI application would result in discrimination or exclusion, if personal autonomy were impaired as a result, or if an AI application would have such positive effects for personal development that a claim to a right to access these applications could arise.

15 Intellectual property

Teun Pouw

Artificial intelligence brings with it fascinating possibilities for application. In fact, not a sector can be thought of that would not be affected to some or large extent or even drastically reformed as a result of this development. Any discussion of such valuable innovations cannot ignore the topic of intellectual property rights. After all, the inventor of an AI system has every interest in the AI system he has invented being protected by the law, or perhaps even more importantly: that the products invented and produced by his AI system enjoy legal protection. But this protection cannot be taken for granted. Important questions are, for instance, whether the particular system and its creations can be protected against use by others. And if so, who is entitled to the intellectual property rights? A robot being an owner will be taking things too far for many people. In this article, I discuss a few relevant aspects of intellectual property (IP) in relation to artificial intelligence (AI).

The background

The term 'intellectual property rights' as a collective name for the rights to creative achievements by a human. These creations are regarded as products of the (human) spirit, the main aspect that distinguishes humans from animals.[1] The starting point that the creations must be the work of a *human* can also be found in various international IP treaties[2] and national statutory IP regulations.

Why do we have these rights? The thinking is usually that innovation should be rewarded. By giving the conceiver or inventor a certain monopoly over what he has created, he is given the opportunity to profit from his creation. He can in that

[1] In the United States, for instance, copyright was denied for a photo taken by an ape: https://cases.justia.com/federal/appellate-courts/ca9/16-15469/16-15469-2018-04-23.pdf?ts=1524502895.

[2] See, for example, the Berner Convention from 1886, to which 176 states are now party, in which 'the author' is the central focus and it emerges from the rest of the text that this term (only) pertains to 'subjects' of the particular states.

case, for instance, charge licence and other fees, which encourages him to invent or discover even more. If this kind of protection is not granted, there would be the danger that others would hijack the creation. The considerable financial and mental efforts of the inventor would not be rewarded in that case, which could have an inhibitory effect on innovation.[3] This basic thinking in IP law does not appear applicable to artificial intelligence in precisely the same way. Artificial intelligence also brings to light the (possible) limitations of the current legal framework. The intellectual property rights that require further discussion in the context of artificial intelligence are, in particular, copyright and patent right.

Copyright

Based on the current legal framework, there are two basic conditions for a copyright to arise:

1) the creation in question must be the work of a human,
2) an adequate degree of creativity must have been involved in the creation.[4]

We can first look at products made with the help of an AI system. The use of artificial intelligence need not stand in the way of copyright arising to the creations in question. For example, it is certainly conceivable that an author may have used an AI system merely as an aid but that the author continues to have sufficient influence on the ultimate result. You could say that a comparison could be drawn to advanced computer systems already in existence and that this is therefore nothing new.[5]

[3] The open-source software movements as well as the creative commons movement take a different view of this.

[4] Or, in the Supreme Court's words: for a work to be created, the work in question must '*have its own, original character and bear the personal stamp of the maker*' (Supreme Court 30 May 2008, LJN BC2153, NJ 2008/556). Subsequent to this, the European Court of Justice also decided in the Infopaq I judgment that what must be involved is '*the intellectual creation of the author of the work*' (CJEU 16 July 2009, no. C-5/08, LJN BJ3749, NJ 2011/288).

[5] In the same sense: Frank Gotzen & Marie-Christine Janssens, 'Uitdagingen van artificiële intelligentie voor het auteursrecht' [Challenges of artificial intelligence for copyright], KU Leuven (2019); Hristov, '*Artificial Intelligence and the Copyright Dilemma*', The IP Law Review, Vol. 57, No. 3 (2017).

The situation will be different if the author uses AI systems that operate more independently,[6] which will be used more and more in the near future. In that case, one will not always be able to say that the author exerted such an influence on the end result that there is still a case of a human role and/or a creative process. The opinion in relation to such creations may in that case have to be that they belong to the public domain and are therefore royalty-free. It does not seem possible in advance to draw a clear line between which AI systems would or would not result in copyright protection for the products produced using them. Such an opinion will ultimately be reserved for the courts, which will have to take into account all the circumstances of the case.

The foregoing can be nicely illustrated with reference to a few examples of already existing intelligent systems that even create artworks themselves. For example DeepDream, Google's own neural network that generates surrealistic landscapes based on an analysis of random photos fed into the computer.[7] Or the one-armed robot e-David, which can autonomously create a painting without basing this on any existing artwork.[8] Another example is an AI system in Japan that wrote a novella and scored well in a literature competition.[9] These examples are, of course, a challenge to the notion that only humans can perform a creative act of creation and pose a challenge for those working in the area of copyright.

For the rest, a closer look can also be taken at the protection of the AI system itself. The system will largely consist of software which, in principle, can be protected on the basis of the current copyright law. In line with the foregoing, it also applies here, however, that a system that independently continues to develop itself will, at a certain point, no longer be able to be qualified as the creative achievement of a human.

---

[6] See also: Ramalho, 'Will robots rule the (artistic) world? A proposed model for the legal status of creations by artificial intelligence systems', Journal of Internet Law (2018).

[7] https://ai.googleblog.com/2015/06/inceptionism-going-deeper-into-neural.html

[8] https://www.theguardian.com/artanddesign/2016/apr/19/robot-art-competition-e-david-cloudpainter-bitpaintr

[9] Frank Gotzen & Marie-Christine Janssens, 'Uitdagingen van artificiële intelligentie voor het auteursrecht', KU Leuven (2019); https://www.smithsonianmag.com/smart-news/ai-written-novella-almost-won-literary-prize-180958577/

Patent right

Patent right deserves a detailed (independent) discussion of its own. For this article, I suffice with the comment that patent right, too, assumes a natural person ('the inventor') and that patent right also requires that there be an intellectual contribution to the inventive process.[10] In patent law, similar topics are therefore relevant as discussed previously in the context of copyright.[11]

Need for perspective?

Earlier I pointed out the justification for the existence of intellectual property rights, specifically, among other things, securing the inventor's livelihood (and in doing so also promoting innovation). It is a legitimate question to ask whether this thinking still holds for AI applications. After all, if a specific person cannot be designated as the inventor, it can also be defended that no one's interests are infringed if another party makes the same product as the AI system.[12] It also emerges from the WIPO's figures that the use of AI systems is taking off significantly.[13] There is evidently enough of a (financial) incentive to use these systems, despite the lack of clarity about their protection.[14] This raises the question of whether (additional) protection is indeed needed for AI systems. In my view, it is nonetheless worth recommending

[10] Among others, Prof. P.H. Blok, *Echte rechten voor kunstmatige creaties Moeten we octrooien blijven verlenen als slimme systemen het uitvindwerk overnemen?* [Real rights for artificial creations Do we need to keep granting patents if smart systems take over inventing?], deLex Amsterdam 2018, page 14.

[11] Patent law also involves its own aspects, for instance the requirement of 'replicability' (which means that a person skilled in the art must be able to replicate the invention, which can prove difficult with AI) as well as the restrictions that apply for obtaining a patent to software. For the rest, the European Patent Office's guidelines confirm that AI-based inventions are indeed patentable if the invention is decidedly technical in nature: Guidelines for Examination in the European Patent Office, November 2018 (https://www.epo.org/law-practice/legal-texts/guidelines.html).

[12] In this sense, Chavannes, who asks whether it is indeed a problem that the IP regulation does not provide full protection for AI systems: Remy Chavannes, 'De bescherming van *deep learning*-systemen door het intellectuele eigendomsrecht' [Protection of deep learning systems by intellectual property law], 5 AMI – Journal for copyright, media and information law (2018).

[13] WIPO (2019), *WIPO Technology Trends 2019: Artificial Intelligence*. Geneva: World Intellectual Property Organization (https://www.wipo.int/edocs/pubdocs/en/wipo_pub_1055.pdf).

[14] In the same sense: Remy Chavannes, 'De bescherming van *deep learning*-systemen door het intellectuele eigendomsrecht' [Protection of deep learning systems by intellectual property law], 5 AMI – Journal for copyright, media and information law (2018).

that new legislation be considered, if only to avoid years of ambiguity – and many lawsuits (as was – and still is – the case for aforementioned earlier major technological developments).

Conclusion

The reality of creations by non-humans raises questions about the legal status of works generated using AI systems and the AI systems themselves. Copyright law and patent law have both survived the rise of computers and the internet (with minimal amendments). A good number of lawsuits were (and are still being) conducted in which these topics are the focus in order to arrive at a correct interpretation within the existing legal framework. The same will occur in relation to artificial intelligence if additional regulations are not introduced. This prompts the question of whether additional regulations are desired. Another view argues in fact for 'letting go' of protection for AI systems. Whatever the case, the development of artificial intelligence forces us to rethink key concepts in intellectual property law. It would also be interesting to conduct further debate on the reasoning on which intellectual property law is based in relation to the (potential) importance of AI for our society.

Considerations

- The current legal framework for intellectual property rights does not seem to provide the desired protection for owners of AI systems. For instance, based on the current regulations, it is not clear whether AI systems and the products created using AI systems enjoy legal protection and whether this can even vary depending on the AI system.

- Without further regulation, a (new) 'lawyers' paradise' could arise. Questions will have to be put to the court on all sorts of aspects relating to AI: who is the right holder, has he/she performed a creative achievement, is there any infringement? Etc.

- One could also ask whether further protection of AI systems is necessary or desirable. This deserves in part a societal approach, in view of the relevant applications that AI can bring with it (in the area of healthcare, for instance).

16 Internet of Things

Victor de Pous

Digitalised objects that interact invisibly with their physical environment are growing strongly in number and relevance because of the quality and efficiency improvements of existing processes and the design of new ones. 'Information at your fingertips' in excellent form. By 2020, the world is expected to have some 30 billion connected 'things' that autonomously function in a network as an aid or means for decision-making. In order to make any sense of the enormous volume of individual and consolidated data, the use of artificial intelligence is unavoidable. From a legal perspective, discussion of the Internet of Things is always accompanied by warnings of (the risks of) security breaches, including personal data breaches. More general questions about AI also arise, for instance in relation to the transparency of the data processing and accountability. But the problems are considerably broader. From access to markets and the use of generated data to the right to software updates and diverse liability issues. In the first place, every process, product or service requires lawful data processing.

Concept

From a general perspective, the Internet of Things is a network of physical devices, machines, vehicles and other unique identifiable objects, which are equipped with embedded electronics and network connectivity, allowing data to be processed in real time. Or to put it more universally: 'a rapidly evolving and expanding collection of diverse technologies that interact with the physical world.'[1] Rumour has it that the term Internet of Things was coined by British computer scientist Kevin Ashton in 1999.[2] Artificial intelligence is increasingly a part of this collection of diverse technologies.

[1] https://www.nist.gov/publications/considerations-managing-internet-things-iot-cybersecurity-and-privacy-risks

[2] http://www.rfidjournal.com/articles/view?4986

The societal and economic interest lies in the fact that everyday 'things', in combination with other information technology like big data and artificial intelligence, primarily in the form of machine learning, are acting autonomously and serving as an aid or means for decision-making, thus improving existing processes and making entirely new ones possible. The Chinese Prime Minister Wen Jiabao presented us with the following sum on 7 August 2009: 'Internet + Internet of Things = Wisdom of the Earth'.

Previous century

The current attention for automated autonomy for decision-making suggests that Internet of Things-like implementations are something new. The first cash dispensing machines in the United States are an example of the phenomenon before the term existed. These date from 1974. Other early applications can be found in networked computer systems, the first generation of GSM telephony, around 1995 (every mobile digital telephone still has a unique 'media access control address' (MAC address) and geolocation).

A 20th-century Dutch example of 'inter-networking things' and at the same time a notable regulation for a concrete application dates from 1989. To prevent overloading the legal system and the ensuing consequences, common, simple traffic violations were removed from the Criminal Code at the time and the Code of Criminal Procedure likewise no longer applied for these offences. The applicable law is – alongside, among others, the Road Traffic Act 1994 – the General Administrative Law Act. A legal trick to enable automation of the administration of justice.[3]

Mulder Act

Via this legal route, a minor speed violation or running a red light could be handled entirely automatically without the involvement of the court, that is to say, from the determination of the violation through to the creation of the dated settlement proposal. (As an aside: the administrative sanction is still sent by old-fashioned

[3] We expressed doubts about this. Does this procedure satisfy the guarantees for due process? In particular, we wonder whether the legislator is 'playing improperly with legal domains' here. See V.A. de Pous, *Recht voor een nationale informatiesamenleving* [Law for a national information society], Amsterdam, 2011. Originally written for the collection *Interoperabel Nederland* published by the Standardisation Forum, part of the Ministry of Home Affairs and Kingdom Relations.

letter, but these days could just as well be delivered to one's inbox on the mijnoverheid.nl website.) In order to make this possible, the so-called Mulder Act made the public prosecutor an administrative body and gave it the right of summary execution. Police officers were also given this power.[4]

A judgment from the Rotterdam criminal court shows how topical the problem of lawful data processing in the Internet of Things is in relation to the 30-year-old Mulder Act. This court ruled on 16 July 2018 that the police were not permitted to use a speed gun on vehicles at the Heinenoord tunnel on the A29 at Barendrecht *because there is a bend in the road there*.[5] Consequently, the radar measurement cannot be correct. It cannot be accurate enough.

Work of man

Although the linked technology does process this type of violation, on grounds of the law only *public officials authorised to do so* may actually impose the sanctions (Article 3(2) WAHV). So for the time being, it is still an actual human who makes the decision. While in a formal legal sense the Internet of Things explicitly serves as an aid in decision-making, the practice makes one suspect otherwise. The information technology might – secretly – be making the decision itself, as follows from a judgment of the Leeuwarden Appeal Court.[6]

In the event of registry control by the RDW (National Vehicle and Driving Licence Registration Authority), the appeal court cannot determine that the administrative sanction has been imposed by a public official authorised to do so. The RDW lists the charging officer code 404040, but it *cannot* be traced from this which investigative official imposed the sanction. The charging officer can only be determined after further information has been provided by the public prosecution office on the procedure for determining whether vehicles are uninsured and the imposing of sanctions for this on registration holders.[7] Apparently, despite it being a fully automated process, the procedures have officially been set up in accordance with the law, which also emerges from a judgment of the Supreme Court.[8]

---

[4] Act of 3 July 1989, for the administrative-law settlement of infringements of specific traffic regulations (Traffic Regulations (Administrative Enforcement) Act -WHAV).

[5] *SUSPECT V. PUBLIC PROSECUTION OFFICE*, Rotterdam District Court, 16 July 2018, ECLI:NL:RBROT:2018:5902

[6] *PUBLIC PROSECUTION OFFICE VS. SUSPECTS*, Leeuwarden Appeal Court, 20 February 2014, ECLI:NL:GHARL:2014:1236

[7] *PUBLIC PROSECUTION OFFICE VS. SUSPECTS*, Leeuwarden Appeal Court, 4 June 2014, ECLI:NL:GHARL:2014:4324

[8] Supreme Court, 16 February 2016, ECLI:NL:HR:2016:240).

From legal transaction to all actions

In the 1980s, it was posited that a Dutch resident was included in, on average, 40 databases. This could be banks and insurers, the Tax and Customs Administration, utility companies, municipalities and water boards, the RDW, housing associations and, for instance, associations, like the ANWB, Veronica and sports clubs. What these systems document and otherwise process are in principle and primarily *legal transactions:* conduct with an envisioned legal effect. Payment, membership, registration in the population register, ownership of a motor vehicle, and more.

If one were to link the various databases of personal data, this would ultimately produce a composite *administrative shadow image* of an individual identified by name; this was the concern expressed 35 years ago by a small group of privacy experts.[9] It is conspicuous that this discussion is being repeated today, *even without the databases being linked*. Thanks to the resounding success of the debit card and contact-free payment, the lion's share of payments in the Netherlands takes place electronically and, whatever the case may be legally, until recently some banks analysed their account holders' transactions or show a commercial interest in doing so. 'Payment data give something of a complete picture of someone's life: what you spend your money on, what associations you are a member of, who you spend time with, what patterns are visible,' the Dutch Data Protection Authority (Dutch DPA) now determines.[10]

The option of an even more complete picture presents itself. Both through the use of the internet and the corresponding large-scale use of 'tracking software', such as tracking cookies, tracking pixels and 'fingerprinting', and Internet of Things applications, all sorts of actual actions are constantly being documented - on an enormous scale - *without any envisioned legal effect of the person*. From the mountain of data, businesses and governments can distil connections and needs[11], enabling them to predict people's future behaviour, without the people even being able to suspect this. As a result of this, organisations can treat some people differently to others. With all the ensuing consequences, according to the Dutch

[9] See in this respect: V.A. de Pous, *Computerrecht* [Computer law], Amsterdam, 1982.

[10] https://autoriteitpersoonsgegevens.nl/nl/nieuws/banken-mogen-betaalgegevens-niet-zomaar-gebruiken-voor-reclame

[11] This information technology is called 'big data'. https://en.wikipedia.org/wiki/Big_data

DPA in 2014.[12] On 13 June 2019, our privacy regulator again warned: 'Internet of Things devices collect data that can say a great deal about you.'[13]

Security

In the meantime, the current focus of politicians and the legislator is mainly on the security aspect of the IoT. The government and European Commission are investigating the possibility of expanding the CE quality mark to include minimum requirements for the digital security of wireless devices.[14] Both the government and the Radiocommunications Agency Netherlands[15] advocate mandatory minimum security requirements for IoT equipment. To date, Europe has not shown much of an interest in this.

On 8 June 2018, the European Council of telecommunications ministers put together a general orientation for the Regulation on ENISA and Cybersecurity certification.[16] In line with the motion from member Paternotte cum sui,[17] the Netherlands at that time advocated for *mandatory certification* for Internet of Things devices, but got no response for this from the other member states. So for the time being, the establishment of a European framework for cybersecurity certification will continue to be on a voluntary basis, with which the European Parliament agreed on 12 March 2019.[18]

Conclusion

Everyone must be able to trust in the lawfulness of the data-processing processes, products and services, regardless of the technology, application or sector and irrespective of the nature of the data. This crucial societal basic condition should apply all the more for the use of the Internet of Things, on account of the black box nature of both the equipment and the self-learning software, the invisibility of the

[12] http://www.cbpweb.nl/Pages/jv_2013.aspx

[13] https://autoriteitpersoonsgegevens.nl/nl/nieuws/ap-geeft-tips-voor-privacy-bij-internet-things-apparaten

[14] https://www.tweedekamer.nl/kamerstukken/kamervragen/detail?id=2018D37188&did=2018D37188

[15] https://www.agentschaptelecom.nl/actueel/nieuws/2018/juni/04/onveilige-iot-apparatuur-risico-voor-samenleving

[16] COM 2017/477, abbreviated as the Cybersecurity Act.

[17] Parliamentary document 21 501-30, no. 422.

[18] http://www.europarl.europa.eu/news/en/press-room/20190307IPR30694/meps-adopt-cybersecurity-act-and-want-eu-to-counter-it-threat-from-china

data processing intertwined with the physical world and the increasing role that the IoT plays throughout society. *This moment of trust has not yet arrived*.
The security of IoT devices is now the focus. As far as personal data are concerned, the mandatory 'appropriate' protection measures are only one aspect of privacy law. The comprehensive processing of personal data in accordance with the law is, in turn, just one part of the standards that the law stipulates for hardware, software and the IoT and the use thereof. Other legal issues are at play as well, among them, perhaps of most importance from the viewpoint of the Dutch economy, free competition.[19] All stakeholders, for a start: user organisations, will not be able to avoid broadening their point of view.

Analyses

- Statutory minimum security requirements for IoT devices could be beneficial, State Secretary Keijzer wrote to the Lower House of Parliament on 20 June 2019.[20] This could ensure that unsafe products are pulled from the market and refused, providers would be given an incentive to comply with legislation and regulation, and citizens and companies would be able to recover damage caused by insecure digital products. The hypothesis touches a sore spot. *Why is the sector not motivated itself to produce safe digital equipment?*

- The processing of enormous volumes of IoT data (big data) requires artificial intelligence to unlock the data relevant for decision-making. The Internet of Things functions as the supplier of the data, which are 'mined' using machine learning and converted into a usable (functional) result via analysis. This applies likewise for data on digital security. In the same way, in the near future, AI-based cybersecurity software will be able to fully automate the prediction and detection of and response to a digital

[19] The monopoly position of big American tech companies is increasingly the subject of discussion and has prompted investigation by regulators. According to the European Commission, Google acted in violation of competition law by, among other things, requiring telephone manufacturers to preinstall the search engine app Google Search and the browser app Google Chrome as a precondition for the granting of a licence for the Play Store, Google's app store.
http://ec.europa.eu/competition/elojade/isef/case_details.cfm?proc_code=1_40099

[20] Parliamentary document 26643, no. 618.

security incident. *This makes it seem as if the use of artificial intelligence, in any event in this perspective, will become a legal requirement.*

- All sorts of factual behaviour can also be monitored and documented without an envisioned legal effect via an Internet of Things application. For this, the data subject must at all times be able to give his permission freely if the application processes his personal data. We argue on grounds of contract law (good faith in the performance of a contract) in combination with privacy legislation *that the full functionality of an Internet of Things object must in principle be available, regardless of whether the user accepts tracking technology.*

17. Auditing: monitoring the black box or black magic?

Maarten Souw

Anyone interested in using artificial intelligence or, for instance, robotics cannot avoid an important preliminary phase: risk assessment. This applies both for government organisations and companies considering using these promising digital technologies in their operations. The assessment emphatically serves as a starting point for the principal and auditor. What is the focus in this process? In every audit, the principal and auditor must ask where the limits delineating the scope of the project lie. In day-to-day practice, the audit budget often proves insufficient to carry out a full audit or to have one carried out. Making choices and honing in on the key risks and control measures are therefore the essence of an audit of this advanced digital technology.

Object

In virtually every audit, the auditor takes stock of risks. Audits can often centre on standard issues, such as the accuracy of financial statements or quality of information security. The issue could, of course, be more specialist, such as the audit of an AI system. In any event, the IT auditor will perform the risk analysis in figure 1.

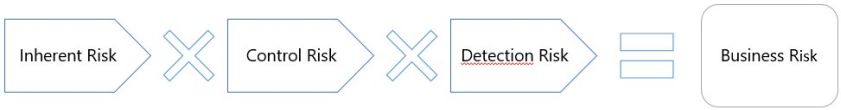

Figure 1 – Audit risk
The risk that when using (AI) technology, a company is dependent on the inherent risk, the security measures in place and the likelihood that a risk (or failure of a measure) has been overlooked

This figure graphically illustrates that every process or every technology entails inherent risks. The principal wants to control these risks and takes measures to this effect. The control measures in turn have their own limitations: a roll cage in a car does not prevent the driver from arriving too late, for instance. The final complication arises if the organisation overlooks deficiencies in the risk

management. This is the significance of the third arrow, the risks that are not provided for or that are poorly managed. These three limitations mean that not all risks are covered. Risk management generally has two axioms:

a) that risks are only relevant if they affect the organisational goals;
b) that covering all risks is not feasible.

Since the principal cannot cover all conceivable risks, he will have to focus on the most important ones. In the event of AI, those are three specific matters: publicity, bias and evolution. These risks are explained in more detail below.

Specific use determines risk

The way in which AI is used determines the risks and, by extension, the monitoring approach: the reason why (need) and the way in which a company uses AI can vary in scope and ambition, after all. The smart robotisation of manual processes is a relatively simple application, like the forwarding of information between two systems. A more complex application of AI would be training an SIEM (alarm system for the digital world) to identify cyber attacks, whereby the system takes over the role of experienced human analysts. These professionals will have a wealth of undocumented knowledge and experience. Modelling this kind of extensive and specialist knowledge is therefore expected to take more effort than automating a simple action.

From risk analysis to audit programme

In the second step, the auditor translates the risk analysis into an audit programme. In this chapter, risks in the *use*, *creation* and *evolution and management* of AI are anticipated. The risks in table 1 apply generally. The thinking and cohesion behind the measures – which must *in principle* be in place – are also universal. Ideally, the measures ensure:

a) a good translation from object to system; and
b) that once the system has been set up, it continues to function.

The measures that the principal still *can* implement depends on the timing of the audit, of course. In the event of a preventative audit, the principal may have more time to take measures. However, all the information may not be available yet. In an audit after the fact, the actual functioning of an AI system can be determined, but the principal will undoubtedly have less room to take measures.

Risk	Control measure	Points for attention in the audit
Legal consequences of the *use* of AI Negative publicity if the *use* of *AI* becomes known.	Organisation has an AI policy: • The use and preconditions for AI have been clearly delineated; • Hum an responsibilities in relation to AI have been clearly laid down.	Does the organisation have a cohesive policy on the use of AI? Has the policy been translated into a cohesive system of controls?
Implicit assumptions have been incorporated in the *creation* of the AI	In the development process, contention was organised, or creative techniques or other techniques were used to prevent tunnel vision.	The auditor investigates how the organisation set up contention. The engagement of an expert or second opinion must be considered on top of that. The functioning of the IT General Controls*) must of course be investigated.
Implicit knowledge has been modelled in the *creation* The AI system *evolves* and steps beyond its design limitations	Set up a monitoring process in which: • The outcomes of the AI system are checked (via spot checks); • (Later) deviations are detected and adjusted if necessary	The auditor investigates whether such a monitoring and adjustment process is in place: • The preconditions and limits of the decisions that an AI system can take must be explicit; • The organisation uses a comprehensive monitoring system.
	Set the system up in such a way that the AI system cannot change system rights.	The functioning of the IT General Controls must have been determined in the audit. The AI system uses known and verified system rights and cannot change its own rights.
*) IT General Controls is a collective term for the key hygiene measures in the information provision. The IT General Controls encompass at least the logical access security, testing, change management and continuity management. The IT General Controls form the basis for the controlled functioning of an information system.		

Table 1. Examples of risks, control measures and points for attention.

Risks and management in use

Examples of risks with the *use* of AI are legal risks and publicity risks. For instance, an organisation can only make decisions on individual persons with human intervention. This provision from the GDPR impacts STP, Straight-Through Processing. Depending on the degree of automated decision-making (STP), it could be wise to seek legal advice. A second risk with the *use* of AI is publicity risk. Various opinion makers (like Bill Gates or Stephen Hawking) have expressed concerns about AI, while the European Union is investigating the risks of AI as well.[1]

A control measure that is the quickest to realise is the use of an *AI policy* in which the management explains its vision and control. Topics include vision, strategy and area of application, limits imposed on the AI systems and how the organisation wants to control the use of AI.

Risks and management in the creation of the AI

The two most striking risks in the *creation* of AI are bias and implicit knowledge. The *bias* issue (prejudice, implicit assumptions) arises because AI systems have usually been based on the existing way of thinking and input of experts. Implicit knowledge was first detected in *rule-based* systems. This manifests because the AI system makes connections that do not exist. Although many apocryphal examples are given, this is a recognised risk.

When assessing both risks, the requirement of the auditor's expertise plays a major role. Auditors are required to only make a pronouncement if they consider themselves sufficiently expert. The auditor must ask himself whether he can recognise this kind of implicit knowledge or bias. At the same time, an interesting dilemma presents itself; after all, the auditor remains responsible for the audit. He will have to investigate the risk of tunnel vision in any event. A solution for this issue can be reached as follows.

First of all, the auditor can make a process-based assessment of the process by which the AI was created. To some extent, this includes the usual IT General Controls, the system of logical access management, testing and change management. The auditor can also investigate whether the principal attempted to prevent tunnel vision, for example by asking the following questions:

[1] See the statements from the Science and Technology Options Assessment of the European Parliament in T. Metzinger, P.J. Bentley, O. Häggström, M. Brundage, *Should we fear Artificial Intelligence*, Euro Parliamentary Research Centre, March 2018 – PE 614.547

- Did the principal investigate what bias (or undocumented knowledge) was incorporated in the AI system?
- Have contentious or creative techniques (such as reversing the goal) been used during the creation of the system?

As a second option, an external expert can be relevant for the substantive assessment. An AI system is, after all, intended to take over the role of a subject matter expert. It is unrealistic to expect the auditor to assess the work of these subject matter experts. It could be the case that the principal has already tackled the uncertainty about proper functioning. The use of an AI system whose functioning has been recognised in the market is a plausible approach. In this case, the auditor can focus on the correct application of the system used. If the principal has not tested the functioning of the AI system, the auditor can consider seeking a second opinion. In both cases, the auditor will have to look into whether the system has been set up properly (do the assumptions in the system tie in with the new area of application?).

Risks and control measures in the evolution and management *of AI*

A typical aspect of AI systems is that they are learning systems. With the usual IT solutions, the builder strives for repeatability and verifiability. AI systems, on the other hand, are intended to improve themselves. It is therefore to be expected that an AI system will change its behaviour over the course of time. An AI system can, while running, take different decisions than it took when delivered. Here, too, the auditor cannot be expected to play the role of an expert. The auditor can contribute ideas on preventing undesired decisions:

- One way is to incorporate hard limits. An organisation can incorporate a gate, for instance, such as a check by a human as the final step. Or an intermediate form could be used, in which only outliers are sent to a human handler. Naturally, an organisation must also take basic measures, such as rights management. An AI system may not simply adjust its access rights the way a natural person would be permitted to;
- A second way is to monitor the results. Monitoring requires a process and the use of an expert. This person must, after all, evaluate the relevant developments. This approach seems laborious, but there is a second argument. The principal expects that the AI system will develop. Figuring out whether an AI system is developing as desired also requires a system of monitoring.

In conclusion

We have assessed the risk management – the auditor's object of inquiry – in relation to the use of AI. The necessary control measures range from straightforward operational measures (hygiene measures in relation to the system) to managerial measures (policy in relation to the use of AI). Which control measures are opportune is the situational choice of the principal. The need for human 'oversight' depends on what impact the automated decisions have, for instance. It is generally acceptable for AI to simplify trivial work duties. On the other end of the spectrum, human oversight for bigger decisions contributes to societal acceptance. The interesting issues arise towards the middle of the spectrum. Here the principal must make a choice as to whether it wants to keep a human handler involved in the process. The auditor can – precisely in preventative audits – add value by indicating the risks clearly with the principal.

At the same time, the traditional audit issues remain relevant. An AI system must be built in a controlled manner, for instance. The auditor will have to monitor the development, change and testing process in this context. Once realised, the AI system must be protected against unintentional changes. This test, too, is part of the traditional assessment of IT general controls.

Points for attention

- For every audit, the creation and management of the AI must be verified.

- The auditor must look into whether the corporate policy (in relation to AI) has been translated into functioning systems. Parallel to this, it must also be determined whether the corporate policy is safeguarded by means of a cohesive system of controls.

- The proper functioning of IT General Controls (standard hygiene measures in IT) remains a basic condition. The IT auditor must investigate whether the IT General Controls are working;

- The correct expertise is essential for the validation of the AI algorithms themselves. The engagement of a second opinion could be considered in this context. The intensity of the audit depends on the ambition level of

the company itself. Above all, the auditor will continuously have to ask himself: '*What are you doing, Dave*?'

18 Liability

Laurens Thissen

Software has for years been providing plenty of food for thought to lawyers endeavouring to answer questions of liability. With Artificial Intelligence (AI), liability issues can cause some serious headaches. After all, AI is a digital technology that – within specific pre-programmed bandwidths – has self-learning capacity and can take autonomous decisions. The legislator did not take these kinds of technologies into account. The question is therefore how liability for damage caused by AI should be assessed. The question also arises of whether application of the liability law in force results in justified or desirable outcomes.

In advance

The fact that the supply of an AI product can result in contractual liability if what has been supplied does not meet the requirements set out in the contract should be clear. After all, the supplier is in that case liable for not performing his obligations under the contract. This article therefore only addresses what is referred to as extracontractual liability, i.e. the liability for AI if no contractual relationship with the injured party exists.

In order to determine which statutory liability regime is eligible for application to AI, AI must first be qualified legally. This immediately presents a problem. After all, there is no clearly delineated definition of what AI is exactly. The technology has a number of variations, all of which have their own methods and sub-methods[1] for being artificially intelligent, and these variations are all still in development. The most common technology today is machine learning.[2] To avoid

[1] A distinction can be made, for instance, between weak and strong AI, responsive or non-responsive AI and AI that has a high level of abstraction, etc. The technical functioning of these falls (far) outside the scope of this article; however, in the context of liability issues, it must be pointed out that given the different technologies, it is possible that no one-size-fits-all liability regime applies or should apply.

[2] Some perspective is needed here, too, since the term 'machine learning' is also used as a collective term for various technologies and sub-technologies for 'an automatic learning capacity'. For instance, machine learning can be achieved by letting self-learning algorithms 'loose' on 'random forests' or neural

overcomplicating the study into liability for AI, only machine learning – in the sense of algorithms with self-learning capacity – is addressed below. This may well be entirely incorrect, because the other AI applications certainly deserve further discussion.

Peculiarities of AI

AI in the aforementioned sense is a technology that analyses data, identifies patterns in data and subsequently draws conclusions from those. By comparing the outcome of an analysis with the outcomes of earlier analyses, it draws increasingly better conclusions. And thus it learns. The conclusions could be trivial, such as 'this is a photo of a cat' or 'you might enjoy "House of Cards"', but could also potentially have a major impact, for instance 'I (here: a self-driving car) do not need to brake' or 'you probably do not have a disease'.

What happens if the conclusion is incorrect? The self-driving car should have braked. Or you were sick after all and should have been treated. Substantial damage ensues. The damage-causing event is the incorrect conclusion. Or better: the actual action (or omission) resulting from that conclusion. But why is the conclusion incorrect? Is this due to 1. the algorithm, 2. what the algorithm itself has learned, or 3. the data that were intrinsically wrong or corrupt? And if the cause is 1, 2 or 3, or a combination of two or more, who is liable for that? And who is liable if the maker of the algorithm is someone other than the creator or supplier of the data? And does it make a difference if the conclusion in and of itself is correct but nonetheless causes damage?[3] And does it matter that the damage could still have been prevented by someone, for example the doctor who applied the disease prediction algorithm?[4]

networks; the oft-used 'deep learning' method is a sub-variety of machine learning which itself acquires learning capacity via different means and different methods, such as supervised, unsupervised, semi-supervised, support and transduction. An understanding of the methods used whereby learning capacity is achieved and therefore on the basis of which - potentially damage-causing - autonomous decisions are taken is, in my view, required in order to determine a justified liability regime. The great variety in machine learning makes inquiry into what a suitable liability regime would be somewhat complex.

[3] Take, for instance, a perfectly programmed AI algorithm which, based on correct data, concludes entirely correctly that the self-driving car should change course and subsequently hit one pedestrian rather than the other.

4 That the doctor who acts on the basis of an incorrect conclusion drawn by an AI algorithm will not, with respect to the patient, be able to hide behind the error in the AI model. The question is, however,

Interesting liability questions thus arise. And as will emerge below, these cannot be automatically answered based on the extracontractual liability law currently in force.[5]

Statutory regulation?

Anyone searching in the Dutch Civil Code for a specific regulation for extracontractual liability for AI (or even just for software) will come up empty-handed. No such regulation exists. There are a few specific regulations for extracontractual liability for - briefly summarised - persons and things.[6] For the rest, the general doctrine of the unlawful act must be relied on.

Of the specific regulations for extracontractual liability, those for product liability and liability for defective items appear to be the most eligible for application to AI products.[7] But both immediately fail as well; they only apply to 'items', which are material objects subject to human control, which does not describe software, algorithms, digital data and (therefore) AI. This is not to say that a product that contains AI is not covered by product liability legislation. However, this legislation does not apply to the AI as such. For instance, a self-driving car will fall under product liability, but the AI itself will not.

Analogous application?

These specific regulations also cannot be automatically applied to AI analogously.

Liability for defective items

The reach of liability for items appears to be too limited for analogous application to AI. After all, it imposes strict liability on the possessor of the item, while – in view of the peculiarities of AI – liability should sometimes, but not always, be allocated

whether in addition to having a claim on the doctor, the patient would also have a direct claim on the party responsible for the AI algorithm.

[5] Some suggest, with a certain degree of seriousness, granting legal personality to intelligent systems as a solution for this. For the time being, this is a legal pipe dream, if only because society is not ready for such a fundamental choice, and because of practical 'inconveniences' it would bring with it, such as the impossibility of (non-programmed) (legal) interaction with such systems and limited means of recovery.

[6] The law has specific regulations for extracontractual liability for children, subordinates, non-subordinates, representatives, defective items, buildings, hazardous substances, landfills, mines, animals and products.

[7] Articles 6:173 and 6:185, respectively, of the Dutch Civil Code.

to the possessor. After all, only the manner in which the possessor uses the AI can be attributed to him – in his capacity as possessor. The quality of the data, the algorithm and what the algorithm itself has learned, in other words the AI in and of itself, is usually outside the possessor's control. Allocation of liability to the possessor is therefore not (always) justified or desirable.

Product liability

With product liability, liability is imposed precisely on the 'producer'. According to the law, that is the party that manufactures the end product, a raw material or a component, or anyone who presents himself as producer by – to put it briefly – affixing his brand or trade name on the product. What is interesting in this respect is that this definition lends itself for analogous application to AI. To put it somewhat simplistically, the self-learning algorithm can, after all, be considered the end product, the data can be considered the raw material and what the algorithm itself has learned can be considered a component.

This kind of analogous application would prevent problems of evidence for the injured party. After all, anyone who qualifies as 'producer' would be jointly and severally liable. If the supplier of the data is a party other than the producer of the algorithm, the injured party can set its sights on whomever is the easier mark. Thus he is not faced with the (impossible?) task of having to demonstrate what the concrete cause of the damage was: the data, the algorithm or what the algorithm itself has learned.[8] To that extent, strict liability for the 'producer' of AI appears desirable, and points of connection can be sought for this approach for the doctrine of product liability.

In the other respects, the product liability regime seems less suitable for analogous application to AI. There are namely a number of possibilities of exculpation for the producer which – in the event of full analogous application of

[8] If the injured party can demonstrate that damage was caused by 'incorrect' data, he can sue the 'party responsible' for the data. If the damage has been caused by an error in the algorithm, he can sue the 'party responsible' for the algorithm. If, however, the damage has been caused by what the algorithm itself has learned, the injured party is faced with a problem, because who is responsible for that? In the absence of strict liability to that end, this kind of damage cannot be automatically attributed to the 'party responsible' for the algorithm since that party has no influence on what the algorithm - which is in and of itself correct - subsequently does or does not do.

these to AI – would result in the producer of the AI never being liable.[9] This thus produces undesirable outcomes. Product liability also assumes the presence of a defect, while the peculiarities of AI can entail that AI causes damage without the AI being inherently defective.

Unlawful act

Since there are no specific statutory regulations for extracontractual liability for AI, the general doctrine of the unlawful act must be relied on.

Liability based on unlawful act is assumed if four requirements have been satisfied. There must be an unlawful act (1) which can be attributed to the perpetrator (2) and there must be damage (3) which bears a causal relationship (4) to the unlawful act. The capacity of the liable person is irrelevant here: it could be the possessor, producer or programmer, etc, of an AI algorithm.

One of the questions is therefore whether there has been an unlawful act: that is a violation of a right, a statutory obligation or a societal standard of care. These criteria can also be easily applied to AI. After all, AI that causes a self-driving car to crash will infringe the right to bodily integrity or property right. AI that excludes minorities from certain positions in public administration will be acting in violation of a statutory obligation not to discriminate. And AI that recklessly takes decisions that cause damage will be acting in violation of the societal requirement of due care.

When demonstrating a causal relationship between the unlawful act and the damage, the injured party can encounter difficulties, however. Someone hit by a self-driving car, for instance, will still be able to demonstrate that if the car had not hit him, he would not have suffered damage. The aforementioned person from a minority group will have a harder time demonstrating that he would have been given the management position if he had not been excluded *a priori* by the AI. The same applies for the patient for whom the AI failed to diagnose an illness: would he have been cured or been cured at lower cost if the AI had reached the correct conclusion?

[9] For example, the possibility of exculpation provided for in Articles 6: 185(1)(b) and (e) of the Dutch Civil Code. On grounds of part b, he is not liable if the defect that caused damage was not present at the time the product was put into circulation. The producer would in that case not be liable for damage caused by what the AI algorithm had learned itself. According to part e, he is not liable if, given the state of the knowledge, it would have been impossible to be aware of the defect, and that will easily be the case for new technologies like AI.

And, last but not least, the injured party is faced with the problem of attribution. After all, the damage is always in a causal relationship with (the decision of) the AI. But nothing can be legally attributed to an 'AI'. After all, an AI is not a person having legal rights.[10] Something can indeed be attributed to the possessor, producer, or user, etc, but without strict liability anchored in the law, (the risk of) the unlawful act of the AI cannot be automatically shifted to them. One could still assert that the action of the AI can be attributed to the possessor, producer, or user, etc, because they have 'set loose' a damage-causing AI on society. In legal terms, this would involve the doctrine of endangerment. Application of this species of doctrine of the unlawful act will not provide the solution desired either. After all, endangerment requires that the possessor, producer or user – to put it briefly – must have been able to foresee the materialisation of the danger, to a certain extent, and could have taken measures to prevent it.[11] Because of the peculiarities of AI described above, that will not easily be the case.

Conclusion

Dutch law does not have a specific regime for extracontractual liability for AI. And the existing regulations for extracontractual liability do not automatically lend themselves for analogous application to AI, because this will not always result in a justified or desirable outcome.

A specific statutory regulation for extracontractual liability for AI, which recognises a certain risk liability for the producer of the AI, therefore seems desirable. How such a regulation should be concretely set up depends on the technology (and methods behind it), but the technology is still developing.

For the time being, liability issues for AI must be assessed against the doctrine of the unlawful act. Strict application of this could result in undesirable outcomes, however. In order to avoid undesirable outcomes, when applying the doctrine of the unlawful act to AI issues, points of departure can be sought in elements of the statutory regulations for extracontractual liability that do lend themselves for analogous application to AI. For issues of attribution, for instance,

10 A person having legal rights is a person or legal entity that can have rights and obligations and who can be subject to legal consequences, such as attribution.

11 These are the basic criteria formulated by the Supreme Court, the so-called 'trapdoor' criteria, against which endangerment must be assessed, namely 1. the degree of probability that the injured party would not see the danger, 2. the likelihood that the danger would materialise, 3. the possible seriousness of the effects of the risk materialising, and 4. the degree of onerousness of the precautions to be taken.

the assessment can take place by analogy to the joint and several strict liability for 'producers' in the sense of product liability.

Insights

- Dutch law does not have any specific regulation for extracontractual liability for damage caused by AI.
- Statutory regulations for extracontractual liability do not lend themselves for comprehensive analogous application to AI.
- The general doctrine of the unlawful act presents the injured party with undesirable problems of evidence and attribution.
- Certain elements in statutory regulations for extracontractual liability lend themselves for analogous application within the doctrine of the unlawful act. This can help prevent undesirable problems of evidence and attribution to some extent.
- A statutory regulation for qualified joint and several strict liability for the producer of AI seems desirable.

19. Conditions for applications in education

Joan Baaijens[1], Leon Dohmen and Liesbeth Ruoff-van Welzen

***'We know more than we can tell.'*[2] Knowledge and expertise, acquired through education, are the unique results of each individual's personal development and, as such, serve as a guide for the course of every individual's life, both in a personal and professional sense. On the one hand, AI increasingly and more frequently offers an alternative to the 'classroom' in the process of knowledge transfer and communication between instructor and student. On the other, AI supports educational institutes and instructors in the decision-making on learning programmes, learning methods and student development. These developments are connected with each other. But we must not forget the risks. Specific attention for the aspects of AI Governance, which are becoming increasingly important as the result of mass digitalisation and AI of and in education, is therefore inevitable. Education is regarded in this article in a general sense, not least because the basic nature of the knowledge transfer, communication and decision-making applies in principle for all forms of education, so from low to high.**

Artificial intelligence

AI refers to the application of advanced data and other technologies based on formal, mathematical techniques including algorithms, which includes automatic learning to interpret events, support and automate decisions and take actions.[3] Applying AI in education enables one to unlock data, information and knowledge on a large scale, across extremely diverse domains, and to process this for the purposes of knowledge transfer and communication between instructor and student. Some examples include:

- distance education such as interactive video lectures;
- simulations;

[1] With thanks to H. J.M. Vermijs, educational technologist, University of Tilburg

[2] Michael Polanyi, The Tacit Dimension, University of Chicago Press, 2009, page 4.

[3] See chapter /../ about AI Governance in this collection.

- visual and audio support for lesson planning and lectures;
- taking tests;
- digital reality workshops and digital assignments;
- collecting and assessing educational and course results;
- programming lessons, lectures and work forms;
- intelligent portal systems;
- digital library;
- experimental measurement and observation techniques;
- education communication using tools such as Blackboard, Canvas, etc.

The use of digitalisation and AI is, among other things, a solution for the increase in scale in education and the efficient use of scarce resources (personnel, instructional materials, buildings, etc). AI and digital instructional materials must support the individual interactions between instructor and student as much as possible, or replace some of these interactions, so that an instructor can serve more students in less time. For instructors, the main concern is to save time; for learners, the quality and flexibility of the subject offered set the standard. AI and digitalisation are, in short, indispensable for education.

Interaction between theory and practice

The (future) application of AI for decision-making support is focused on education assistance, learning analytics, personalisation of education and testing.[4] Developing new knowledge and insights is accompanied by the unlocking of diverse large data sets (big data). In this context, it is worthwhile to point out the difference between practical/operational knowledge and abstract/theoretical knowledge. The first involves mainly procedures and instructions on how to act in practice. The second form concerns more reflection on and understanding of practice, with a main focus on finding possible explanations for various issues. However, the transfer of knowledge and expertise is the focus in education for both forms of knowledge.

The development of new knowledge often arises as a result of further exploration of existing knowledge and practical applications, among other ways through combining experimentation and trial and error practices with theoretical

[4] https://www.rijksoverheid.nl/documenten/rapporten/2019/02/28/kunstmatige-intelligentie-ai-in-het-onderwijs

insights and empirical (or, as the case may be, applied) research.[5] AI is now very important for both applying existing knowledge and developing new knowledge. In particular combining existing, different and large datasets often results in the discovery of new applications and also in the encouragement of new knowledge areas. This last effect occurs often because new questions are raised.[6]

Privacy protection

We started with the point of departure that the knowledge and expertise of each individual are ultimately a strictly personal quality. The development of this quality, by following education, therefore always has a strictly personal side. This means, among other things, that the interactions between instructor and student are crucial for the success of the education process. As such, the protection of personal data is constantly at issue in education and, as a result, is an important task. The dilemma that this creates is that the instructor wants to facilitate the student as much as possible and provide him or her with access to everything that entails. On the other hand, the individual student's personal data must remain secure. Examination and test results are strictly personal, for instance.

As such, all information transactions in education are - to a certain extent - bound by laws and information and other regulations. These laws and regulations aim to protect the student's personal life and at the same time make it possible for the instructor to work effectively. This working relationship is the result of a balance between the different respective interests of the instructor and the student. This is why various campuses of Dutch educational institutions host projects centring on topics such as: 'Dealing with each other in the digital world'. After all, digital cooperation requires proper and safe manners. This safe learning and working environment implies specific responsibilities for everyone involved. Because of divergent interests – for instance quality, speed, privacy protection, etc – these responsibilities can clash with each other.

[5] What is interesting in this context is that 'engineering' often refers to a specific combination of practical and theoretical insights. In relation to this, see W.G. Vincenti, *What Engineers know and How they know it*., London, 1990

[6] For example, see developments in 'elementary particle physics' and the origin of big data applications in 'astrophysics'.

Risks for human rights

Protecting the student's personal life, preventing bias in the decision-making and efficiently serving large numbers of education users on a large scale results in an inevitable mixing of moral-ethical and effective-business decisions. The use of AI (and digitalisation) in education puts further emphasis on this. This combination of moral-ethical and effective-business decisions implies a certain cohesion among three types of knowledge. Put briefly, it concerns:

- substantive knowledge (substantive-material matters);
- logical-mathematical knowledge (formal logical structure and models of mechanisms, working in information sets) and;
- moral-ethical knowledge (standards and values more or less solidified in rules).

Particularly this last point – often in the form of accepted codes of conduct – is the weakest link of these three. After all, these rules can be neglected or entirely ignored. It is also possible that the desired standards of behaviour have not yet been put in operation and that they still need to be created through amassed experience and improvisation.

That is why, with a view to AI as a support in the decision-making, it is extremely important in education processes to have insight into the underlying algorithms that simulate the desired basic mechanisms. In the event of an unclear or latent decision-making structure of composite algorithms, there is a very real risk of unverifiable considerations based on unclear criteria. These are the risks that all participants and persons involved are exposed to in case of an uncontrolled introduction of AI in education. Students therefore have the right to know to what extent AI is being applied in education and to what degree this affects decision-making regarding their learning development.

In all of this, for a good understanding of the policy formation in education, it is interesting to make a distinction between, on the one hand, professional administrators and managers and, on the other, direct educators (those involved as instructor or otherwise directly engaged in the primary process of education). Administrators and managers assess facilities and AI support mainly in terms of effectiveness and efficiency while instructors (and others directly engaged) primarily want to deliver a substantively good educational product. This generally inherent opposition can, as a result of a mixing of motives, jeopardise the protection of individual human rights of students.

'Humanities' versus 'sciences'

The use of AI in education is part of broad and continuing digital support for an organised form of knowledge and expertise transfer in our educational institutions. In this context, technically feasible possibilities and institutionally embedded standards and values are battling for priority in our current education system.

It is perhaps illuminating in this context to point to British writer and scientist C.P. Snow in his pioneering essay on the existence of two intellectual cultures.[7] He asserts that our society, with the education system contained within it as a large intellectual system, is characterised by, among other things, two fundamental knowledge cultures. The arts, philosophy and the humanities on the one hand, and the fundamental and technical exact sciences on the other. Because of this dichotomy, we see, due to the tumultuous developments in both fields of knowledge and underlying disciplines, that the development of digitalisation (including AI) and the establishment of a humane society (respect for personal lives in an 'open' society) are increasingly difficult to reconcile with each other in advance. Added to this is the fact that many administrators and policy-makers with a humanities background often have to decide on science and technology, while they are not substantively trained in the knowledge areas of the sciences.

Conclusion

In order to protect individual human rights (including in the aforementioned situation of bias in AI-based decision-making) and students' personal lives in education, the two 'knowledge cultures' need to structurally and effectively communicate with each other. That is why AI and digital facilities in education need to constantly be critically monitored, with constant assessment as to whether adequate (ethical) rules and criteria are being used in applying and using AI and digital facilities. After all, the developments of AI (and digitalisation) in education are too important and valuable to leave exclusively to IT and AI professionals.

Considerations

Given the foregoing, the following three considerations are important in order to adequately identify the relationship between the effective and efficient application of AI in education and the protection of individual human rights:

[7] The two Cultures, C.P. Snow, Cambridge 1959, 1993

1. Digitalisation and advanced forms of knowledge access and knowledge transfer with AI are necessary in order to keep the quality of the education system on level despite the exponential growth in scale.
2. Because of the continuing digitalisation and application of AI in the educational process, the personal lives (including privacy) of students must increasingly be better protected. Adequate AI Governance provides a solution here.
3. Students have the right to know how AI is being applied in education and how it affects decision-making concerning their learning development.

20. Towards statutory design standards?

Victor de Pous

De Nederlandsche Bank argues for sector-specific requirements for the use of artificial intelligence by financial companies which require constant monitoring during the use of this transformative technology and are principally decisive in the design phase. With English-language guidelines for the time being, DNB wants to engage in discussion with the market about the possibilities and risks of financial AI. It is not clear as yet whether the debate will result in self-regulation or new legislation and regulations. What is certain, however, is that a regulation trend that has been about for some time is gaining ground in the digital domain: special statutory basic conditions for the design of processes, products and services. While a comparison of DNB's proposals to the European Commission's generic framework evidences differences in the requirements stipulated and in the definitions used, a similar approach does prevail.

Transport electricity

Because of the power outage in Diemen on 27 March 2015, approximately one million households in the provinces of North Holland and Flevoland were without electricity for one hour, while train travel and air traffic were disrupted, hospitals had to cancel surgeries and telephone service dropped. The responsible regulator, the Netherlands Authority for Consumers & Markets (ACM), had an investigation carried out. Study number one revealed a combination of defective technology and human error.[8] A second and broader investigation, this time by the regulator itself, showed that national grid manager TenneT had violated the statutory standard for a 'simple outage reserve' due to how the Diemen station had been designed,[9] a

[8] https://www.acm.nl/nl/publicaties/publicatie/16370/Onderzoek-ACM-naar-stroomstoring-Diemen-afgerond

[9] https://www.acm.nl/nl/publicaties/tennet-voldeed-bij-stroomstoring-diemen-niet-aan-wettelijke-verplichtingen

conclusion that the administrative court later rejected.[10]

What we are concerned with is the statutory standard. This states that an outage in part of the grid, for example a circuit or transformer, may *not* result in disruption of transport, according to Article 31(12) of the Electricity Act 1998. Call it 'operational continuity by design' and an example of a concrete sectoral (*vertical*[11]) quality standard safeguarded in legislation with the aim of minimising to the extent possible the likelihood of an incident, in this case outage of the power supply.

Personal data processing

In the General Data Protection Regulation (GDPR) we also see this preventative line in the *horizontal* design standards stipulated in relation to data-processing operations, so irrespective of technology, sector or application. Article 25 GDPR requires the controller to ensure data protection by design, including through standard settings (privacy by design and privacy by default), via the taking of 'appropriate technical and organisational measures'. In a technical manner, the law forces careful and responsible treatment of personal data by assuming *data minimalisation* (processing of as few personal data as possible), *purpose limitation* (personal data being collected for a specific purpose may not suddenly be used for an entirely different purpose) and *privacy by design* measures (for example pseudonymisation[12]) from the design phase.

The practice shows that producers do not always take this into consideration. A Dutch telemetry case illustrates this.[13] After the Dutch DPA ascertained in 2017 that Windows 10 Home and Pro were in violation of the GDPR with these measurement and communication functionalities, it emerged a year later – based on a data protection impact assessment[14] performed externally – that

[10] *TENNET TSO BV V. ACM AND TATA STEEL IJMUIDEN BV*, Trade and Industry Appeals Tribunal, 12 March 2019, ECLI:NL:CBB:2019:98

[11] 'Vertical' refers here to rules that only apply for a particular branch of industry, as opposed to 'horizontal', which classification indicates that the rules are applicable for every individual organisation, irrespective of sector or type.

[12] For the definition, see Article 4(5) GDPR.

[13] Regarding this technology, see the Dutch DPA's blog. https://autoriteitpersoonsgegevens.nl/nl/nieuws/techblog-telemetrie-windows-10

[14] On grounds of Article 35 GDPR.

Microsoft software posed a 'high privacy risk' for 350,000 government officials.[15] An improvement plan was supposed to provide a solution. The recent updates to the software (together with contractual and organisational changes) were apparently not enough, because the government deemed supplementary privacy agreements with this supplier necessary.[16]

Financial service provision

Just as in other branches of industry, the use of artificial intelligence by financial companies is growing and in this sector, too, the aim is the anticipated improvement of the operational process, product and service thanks to smart systems that to a certain extent act autonomously, on the basis of which decision-making takes place. De Nederlandsche Bank removes any doubt as to whether AI is purely a matter for the future. Existing applications focus on verifying the identity of customers, analysing transactional data, tracking down fraud in claims handling, pricing in bond trading, the automatic analysis of legal documentation, customer account management and, for instance, risk management; an impressive list.

DNB puts forward proposals in the form of six guidelines to limit the damage that an AI-related incident could cause for the financial institution, its customers and ultimately the financial system. In summary: AI must be used exclusively in a *responsible way* at financial service providers, *based on controlled and ethical business operations*.

SAFEST principles

According to the regulator, this fundamental central starting point means that *in the development* of applications, businesses must take into account aspects such as soundness, accountability, fairness, ethics, skills and transparency (SAFEST), according to the guidelines published on 25 July 2019.[17] These aspects constitute a framework in which financial businesses can shape the use of AI in a responsible manner.

However, DNB is not the only formidable party researching and formulating ethical and legal standards and design standards for artificial intelligence. Among

[15] https://www.rijksoverheid.nl/documenten/rapporten/2018/11/07/data-protection-impact-assessment-op-microsoft-office

[16] https://www.rijksoverheid.nl/actueel/nieuws/2019/07/01/overheid-krijgt-meer-grip-op-microsoft

[17] https://europa.eu/rapid/press-release_IP-19-1893_nl.htm

others, the OECD[18], European Commission[19] and, for instance, the G20[20] are getting involved in this domain. To get an idea of a possible shared policy direction, the proposals from the Dutch regulator have been placed alongside the corresponding principles of the European Commission below.

European Commission – 7 guidelines for the design of AI systems in general	DNB – 6 guidelines for the design of AI systems in financial service provision
Human agency and oversight: AI systems must enable just societies through humans exerting influence and fundamental rights being supported. The autonomy of the human may not be diminished, curtailed or abused.	• **Justice**: *fairness* is vital for society's trust in the financial sector: AI applications must not inadvertently disadvantage certain groups of customers. (3)
Robustness and safety: Reliable AI needs algorithms that are so safe, reliable and robust that errors and inconsistencies can be tracked down and tackled throughout the life of the AI systems.	**Soundness:** AI applications must be reliable and accurate, behave predictably, and operate within the boundaries of applicable rules and regulations. (1)
Privacy and data governance: Citizens must have full control over their own data and data relating to them must not be used to damage or discriminate against them.	
Transparency: The traceability of AI systems must be guaranteed.	**Transparency:** *Transparency* means that financial firms should be able to explain how and why they use AI in their business processes and how these applications function exactly. (6)
Diversity, non-discrimination and fairness: AI systems must take into account the entire range of human capacities, skills and requirements, and they must be accessible.	**Ethics**: because AI applications take on tasks that previously required human intelligence, *ethics* also plays a role: financial firms must ensure that their customers and other stakeholders are treated properly and do not

[18] OECD Council Recommendation on Artificial Intelligence.
https://legalinstruments.oecd.org/en/instruments/OECD-LEGAL-0449

[19] http://europa.eu/rapid/press-release_IP-19-1893_nl.htm

[20] https://g20trade-digital.go.jp/dl/Ministerial_Statement_on_Trade_and_Digital_Economy.pdf

	suffer any damage as the result of the use of AI. (4)
Societal and environmental well-being: AI systems must be used to promote positive social changes and increase sustainability and ecological responsibility.	
Accountability: Mechanisms should be put in place to ensure responsibility and accountability for AI systems and their outcomes.	**Responsibility**: businesses must be prepared to assume their *responsibility*: if AI applications unexpectedly do not function properly, this can cause damage for the business, its customers and other important stakeholders. (2)
	Skills: Within a financial undertaking, every individual - from the shop floor through to the board room - must have the appropriate level of *skills* and be familiar with the advantages and limitations of the AI systems with which they work. (5)

Table: the seven guidelines from the European Commission (in development) for the design of AI systems in a general comparison with the six guidelines published by De Nederlandsche Bank (draft) for the design of AI systems for financial services. The numbers in brackets refer to the order used by DNB.

Both frameworks for AI system design exhibit the same tenor: lawfulness, reliability, transparency, explicit responsibilities, non-discriminatory in nature and other ethical principles. In short, principles that put humans at the centre. Internationally, the term *human-centred/centric artificial intelligence* has come into fashion.[21] That could hardly be any different. We also point to differences in definitions between the EU and DNB, for instance in relation to transparency. DNB assumes a larger scope and adds an obligation to justify why AI is even being used.

In addition, two isolated requirements stand out from each other. While the European Commission opts for 'full control' for citizens over their data (a principle that, broadly speaking, follows from the European privacy legislation) and evidently

[21] '*Human-Centered Artificial Intelligence*' is, according to MIT, 'the design, development, and deployment of *artificial intelligence* systems that learn from and collaborate with *humans* in a deep, meaningful way.' https://hcai.mit.edu/ The Dutch Digitalisation Strategy 2.0 talks about '*human-focused AI by design*'.

only advocates human-centred AI that has *positive societal effects*, DNB focuses attention on an entirely different aspect: the *right level of expertise* for literally everyone in the financial company, from high to low, who works with AI systems. This means not only during the design, but also during the production phase.

Conclusion

Technical knowledge has traditionally been the most relevant condition sine qua non for the design of *electronic data processing systems*, but this leitmotif has broadened over time to include organisational expertise. At a later point, those involved in the design phase of digital processes, products and services could no longer avoid having to have *a certain understanding of the legal ground rules* for the technology itself (such as software law) and the concrete application. Today, the urgency of *specific legal expertise* in relation to this is rapidly increasing. The GDPR alone unequivocally evidences this. After all, 'privacy by design' is a must for every producer of a (new) process, product or service that involves the processing of personal data.[22]

De Nederlandsche Bank has now started the debate on *sector-specific* regulations for 'responsible artificial intelligence by design'. In order to minimise the risks of incidents and likelihood of damage, AI must exclusively be used in a *responsible manner* in financial service provision, *on the basis of controlled and ethical business operations*. With this, DNB makes a move with the potential to reach beyond its own sector.

Analyses

- In addition to technical, organisational, ethical and legal principles, the importance of the *societal aspects*, both in the design and use of ICT, is increasing, not least when it comes to a transformative technology like artificial intelligence.

[22] The measures for network and information security made compulsory on grounds of various regulations can also – asserted cautiously or seen indirectly – be regarded as design regulations (security by design). And what to make of the impending European regulation (COM(2018) 286 final) for the design of new vehicle models with mandatory digital functionalities and systems for 'road safety by design'?

- This means that organisations must have – internally and/or externally – sufficient workers with various levels of diverse expertise. The bar is (once again) being raised in terms of competencies. *Can education and the labour market in the Netherlands respond to this demand both in terms of quality and quantity?*

- In our view, De Nederlandsche Bank correctly links the degree to which the use of AI is important in a financial undertaking's decision-making to the potential consequences it could have for the business and its customers for a (i) responsible and (ii) justifiable use of AI. *If the importance of AI in the decision-making increases, the requirements increase as well* (as does the supervision).[23]

[23] This line is familiar from the privacy legislation: the more sensitive in nature the personal data are, the more onerous the requirements stipulated for security.

List of authors

- Dr. J. Baaijens is a Senior Research Fellow at the Department of Organisational Sciences at the University of Tilburg.
- L. Dohmen RI works as a consultant at KEMBIT and is chair of the KNVI Special Interest Group Research and Education.
- N.H.A. van Duuren is attorney-at-law partner for IT, IP & Privacy at De Clercq Lawyers and Notary in Leiden/The Hague and chair of the KNVI Special Interest Group IT and Law.
- J. van Helden is attorney-at-law for IT, IP & Privacy at De Clercq Lawyers and Notary in Leiden/The Hague and member of the KNVI Special Interest Group IT and Law.
- Dr. K. van Ierlant is partner in Data, A.I. & Digitalisation at Verdonk Klooster & Associates.
- W.M.J. Kuik is a manager at Trifork, a consultant at Kweekel Advies & Training and a lecturer at Utrecht University of Applied Sciences.
- W.P. Kweekel RI is director of Kweekel Advies & Training. He is also a management consultant, management trainer and author.
- B. Mulder is an eSkills consultant at De Staffing Groep and stakeholder in AI and ethics at Project SHERPA.
- L. Neddermeijer is founder and director of product strategy at Nétive VMS in Rijswijk.
- Dr. S. van Otterloo is a consultant on software development, privacy and AI at the ICT Institute.
- B.C.P. Pieters Pieters DMS MBA is entrepreneur and director at HB Branding.
- V.A. de Pous is an independent IT lawyer and analyst in Amsterdam and co-founder and board member of the KNVI Special Interest Group IT and Law.
- A.P. Pouw is attorney-at-law for IT, IP & Privacy at De Clercq Lawyers and Notary in Leiden/The Hague.
- L. Ruoff-Van Welzen is chair of the KNVI Special Interest Group Digital Skills, member of CEN TC 428 IT Professionalism and Digital Skills, Director IP3 (IFIP) and director of LRWA in Voorburg.

- Dr. R. Sanders is attorney-at-law for disciplinary and complaints law at De Clercq Lawyers and Notary in Leiden/The Hague and chair of the Disciplinary Tribunal for Occupational Consultants.
- H.L. Souw RE CIPP/E is Security Officer at Nationale Nederlanden in The Hague and chair of the KNVI Special Interest Group IT Audit and Risk.
- L.L.A.M. Thissen is attorney-at-law for e-Tech, IE, Privacy and Media at T-IP Legal in Heerlen and Eindhoven and is secretary of the KNVI Special Interest Group IT and Law.
- Dr. D.G. Velev is professor in computer science at the Department of Information Technology and Communications of the University of National and World Economy in Sofia (Bulgaria).
- Dr. D. de Wit CMC is adviser at O&i Management Consultants in Zeist, chair of the KNVI Special Interest Group eHealth and involved in the Digital Skills in the Healthcare Sector coalition.
- Dr. P. Zlateva is associate professor in system engineering at the Bulgarian Academy of Sciences in Sofia (Bulgaria).

Most of what people have heard about artificial intelligence (AI) probably has to do with systems that can beat humans at a game, such as chess, a television quiz or poker; today there are serious business applications involved. The financial sector is a fine example. Financial service providers are already using AI, for instance to verify the identity of customers, analyse transactional data, track down fraud in claims handling, set pricing in bond trading, and more. The smart assistants provided by Amazon, Apple, Google and Microsoft are also striking examples. We can no longer dismiss AI as something in the far-off future; it must be recognised as a special and 'transformative' technology. 'Of all new digital technologies, artificial intelligence is expected to have the greatest impact on the economy and society in the coming decade,' according to the Dutch government, which has now made it a national priority.

The general sentiment is that AI has enormous potential to improve our lives. Reference is often made to unprecedented results in medical research, preventing financial crises and, for example, the self-driving car, which can improve both traffic safety and flow. This possibility presents itself because intelligent systems can analyse their surroundings and then take action, to a certain extent independently, to realise specific objectives, according to the European Commission's description of AI. This aspect forms the crux of the matter: the use of self-learning information technology as an aid or means for automatic decision-making.

The flipside to this is obvious. In the Netherlands, citizens are concerned about discrimination and privacy breaches, but also about inaccurate data possibly resulting in incorrect decisions and decisions taken by AI systems being generally difficult to reverse. This fear seems to be well founded. There is a threat to democracy, the rule of law and our freedoms, and the potentially irreversible nature of the effects that the smart technology could cause. Alongside the dominant sentiment on 'the silent power of AI' (improving our welfare), the - draft - ethical and legal standards frameworks for human-centred AI therefore also have a complementary common thread: safeguarding our welfare.

AI is a perfect example of a topic that requires a multidisciplinary approach. After all, it is important to take a broad approach to the countless issues and challenges related to - the use of - AI in our society.

ISBN 978-90-9034111-8
NUR 820

knvi