Prescientific Paper

Opportunities and Concerns of Artificial Intelligence with Special Regard to Computer Vision

submitted by

Krystof Hanslik 6th A 2018/19

Supervisors: Mag. Hannes Rechberger & Ostr. Prof. Susanne Reif-Breitwieser

Prague, April 2019

Abstract

Artificial intelligence can help computers solve problems that could not have been solved before. Systems incorporating computer vision can be used in the field of medicine as well as self-driving vehicles.

In medicine, computer vision can help doctors make informed decisions by analyzing medical images. It may be able to replace doctors in parts of the world where healthcare is not widely accessible.

In the car industry, computer vision powers self-driving vehicles, some of which drive on public roads already. With improved accuracy and precision, these vehicles could lower the death toll on roads.

Artificial intelligence also poses danger to liberty and democracy. When used by governments, surveillance can infringe on users' privacy. That is especially disconcerting in authoritarian regimes. Likewise, private companies can make use of this technology to collect information about individuals. In the latter case, regulation such as the General Data Protection Regulation in the EU protect users from this kind of surveillance.

Artificial intelligence will cause people to lose their jobs. This has not only an economic effect, but also raises questions regarding ethical principles that guide robots when working.

Abstract

Künstliche Intelligenz ermöglicht es, dass Computer Probleme lösen, die sie zuvor nicht lösen konnten. Systeme, die computergestütztes Sehen einsetzen, bieten Vorteile, die sowohl im medizinischen Bereich, als auch in selbstfahrenden Autos genutzt werden.

In der Medizin kann computergestütztes Sehen durch die Analyse von Bildern Ärzten helfen, richtige Entscheidungen zu treffen. In Ländern, wo ärztliche Hilfe nicht zugänglich ist, könnte diese durch künstliche Intelligenz ersetzt werden.

In der Autoindustrie wird computergestütztes Sehen für selbstfahrende Autos eingesetzt. Manche von ihnen fahren schon heute auf öffentlichen Straßen. Da sie genauer und präziser fahren können als Menschen, könnte durch ihren Einsatz die Anzahl der Todesopfer auf Straßen verringert werden.

Künstliche Intelligenz kann aber auch eine Gefahr für Freiheit und Demokratie darstellen. Die Überwachung durch die Regierung verletzt die Privatsphäre von Bürgerinnen und Bürgern. Dies ist insbesondere in autoritären Regimen beunruhigend. Auch private Firmen können künstliche Intelligenz einsetzen, um Informationen über Benutzerinnen und Benutzer zu sammeln. Im letzteren Fall schützen Verordnungen wie die Datenschutz-Grundverordnung (GDPR) in der EU vor einer solchen Überwachung.

Künstliche Intelligenz wird Verluste von Arbeitsplätzen verursachen. Dies hat nicht nur ökonomische Auswirkungen. Es stellen sich auch Fragen zu ethischen Prinzipien, die Roboter in ihrer Arbeit beachten sollten.

TABLE OF CONTENTS

1 Introduction

The topic of this prescientific paper has been inspired by countless futurists' predictions regarding technology. AI is often deemed to be the answer to many of the problems that we face today. It is supposed to reduce the frequency of car crashes, provide us with better healthcare and generally improve our quality of living. However, it can also advance oppression in totalitarian states and infringe on our right to privacy. Keeping this in mind, I divided this paper into three sections. The first section explains various terms. The second part analyzes potential benefits, focusing specifically on computer vision. The third part deals with possible dangers of AI as a whole.

1.1 Description of Research Methods

I have researched and compared various sources, mainly in relation to their geographical origin. The research was mostly qualitative.[1] When comparing the American and the European approach to AI, I have determined that the best way to do this was to contrast law proposals in the US with existing regulation in the European Union. My main sources are studies and other scientific sources. The evidence provided is theoretical as well as empirical.[2]

[1] (cf. Adams et al., 2014)

[2] (cf. Adams et al., 2014)

2 Areas of Artificial Intelligence

2.1 Artificial Intelligence

Artificial intelligence (AI) is the ability of a computer "to perform tasks commonly associated with intelligent beings".[3] The exact definition of artificial intelligence (AI) constantly changes. This can be explained by the "AI effect": AI is often considered to be "whatever has not been done yet". This phenomenon has been acknowledged as early as 2002, when Rodney Brooks, the former director of MIT's Artificial Intelligence Laboratory, said, "Every time we figure out a piece of it, it stops being magical; we say, 'Oh, that's just a computation.'"[4]

A different approach to characterizing AI is to compare it to human intelligence. General intelligence as it is displayed by the human brain is currently unrivaled.[5] However, the brain has been surpassed by AI in certain highly specific competences. One area where AI has publicly proven its superiority is the board game Go, when, in 2016, AlphaGo, an artificial intelligence trained for this exact activity, has beaten world champion Lee Se-dol. This is an extraordinary feat because it requires a sort of intuition that was not present in previous challenges.[6]

Sir Tim Berners-Lee, the inventor of the World Wide Web, describes AI as follows: "AI is something that compiles the truth. It's something that can pretend to be a human or can convince a human being that it's a human."[7]

The term "artificial intelligence" is an umbrella term encompassing different fields and methods, including the ones described below.

[3] (Hosch, 2009a)

[4] (Kahn, 2002)

[5] (cf. Villani, 2018a)

[6] (cf. Paul-Choudhury, 2016)

[7] (Campaign Asia-Pacific, 2015)

2.2 Machine Learning

Machine learning is a discipline concerned with the implementation of computer software that can learn autonomously.[8] It is currently one of the most popular areas in AI.[9] A machine learning technique consists of the learning phase and the inference phase. The learning phase uses input data to find the parameters necessary for the task. In the inference phase, the task is performed.[10]

2.2.1 Methods of Learning

In supervised learning, an algorithm is developed by learning from a large sample of pre-labeled data. This could be learning to recognize cats, for instance, by “looking” at a million pictures of cats. The program would then be able to recognize a cat on a picture it has not seen before.[11]

Unsupervised learning is similar to its supervised counterpart, the difference being the lack of labels. The computer finds the underlying pattern itself.[12]

In reinforcement learning, the machine experiments with the environment and learns by reacting to a “reward”. This is how AlphaGo (mentioned above) learned to play the game of Go.

2.3 Computer Vision

Computer vision focuses on understanding images and videos. It is absolutely indispensable in fields such as driverless cars and in clinical settings. Ultimately, its goal is to emulate the human visual system, even surpassing it.[13] Recently transformed by a new approach enabled by faster GPUs, computers are already showing better results than their human counterparts in certain classification tasks.[14]

[8] (cf. Hosch, 2009b)

[9] (cf. Stone et al., 2016)

[10] (cf. Villani, 2018a)

[11] (cf. Villani, 2018a)

[12] (cf. Waymo, 2017)

[13] (cf. Huang, 1996)

[14] (cf. Stone et al., 2016)

> *“Computer vision is the science and technology of machines that see, where seeing in this case means that the machine is able to extract information from an image that is necessary to solve some task.”*[15]

2.3.1 Image Classification

Image classification is a discipline in the field of computer vision concerned with classifying an image into a category, i.e. saying what can be seen in the picture. Image classification is limited to one object per image.[16]

2.3.2 Object Detection

Object detection is a field of computer vision dealing with the recognition of objects in images and videos. Unlike image classification, object detection is able to detect multiple objects in an image and highlighting them with so-called bounding boxes.[17]

Object detection is much more complex than image classification, since the number of objects in a scene is not known beforehand, as well as the size of the objects. Historically, these problems have been solved using approaches such as “sliding windows”. This, in essence means cropping the image multiple times and running an image classification algorithm on all of these crops. The size issue was tackled using variable-sized sliding windows, making these algorithms extremely inefficient.

A framework proposed in 2001 was the first one to be used widely, for example in point-and-shoot cameras that highlighted faces in a picture to focus on them.[18]

After deep learning started to become increasingly popular and capable, a new method called R-CNN (“Regions with CNN features”) was used. It offered many improvements over non-deep learning methods but was plagued by its own problems like difficult training.[19]

[15] (Yoshida, 2011)

[16] (cf. Sánchez et al., 2013, p. 222)

[17] (cf. Felzenszwalb et al., 2010)

[18] (cf. Jones and Viola, n.d.)

[19] (cf. Girshick et al., 2013)

While legacy object detection algorithms worked by attempting to classify multiple parts of an image, a newer system called YOLO, or “You Only Look Once”, only looks at an image once. Using a single neural network enables greater speed which is crucially important in use-cases such as autonomous vehicles.[20] Capable of classifying objects in real time, with high frame rates, YOLO’s speed becomes acceptable for deployment in self-driving cars as well as other applications, as it is meant to be deployed in a variety of fields.

[20] (cf. Redmon et al., 2016)

3 Possible Opportunities in Applying Computer Vision

3.1 Computer Vision in the Field of Medicine

Although artificial intelligence and computer vision are not yet widely used by doctors, this is slowly changing. A new AI and cloud-based system intended to be used on cardiac patients has received clearance by the American Food and Drug Administration in January 2018, signaling that a change might well be under way.[21]

3.1.1 Training

The issue of finding data that can be used during the learning phase of supervised learning (see chapter 2.2) is virtually non-existent in this field, as most medical images are made electronically. Together with data from an EHR (electronic health record) system, an algorithm could learn from large samples of data. DeepMind, an AI lab that has been acquired by Google in 2014[22], has already been granted access to 1.6 million patients' data from three hospitals.[23] Even so, hospitals have only started to store patients' records electronically in the past decade. Moreover, analysis of images in healthcare requires not only parts to be labeled, but also needs their severity interpreted by doctors. Training the model is therefore not as straightforward as it might seem at first glance.[24]

3.1.2 Use Cases

Computer vision in the medical field can mainly be used in the following applications: clinical decision support, patient monitoring and devices to assist in surgery or patient care.[25]

Computer vision specifically opens up a new way of analyzing data and deducting conclusions that human researchers were not able to do. By taking large samples of

[21] (cf. Bazzoli, 2018)

[22] (cf. "About Us," n.d.)

[23] (cf. Quinn, 2016)

[24] (cf. Stone et al., 2016, p. 27)

[25] (cf. Stone et al., 2016, p. 25)

data, underlying patterns and associations can be found. For example, it is possible to predict cardiovascular risk factors by analyzing photographs of patients' retinal fundus.[26]

Another instance of computer vision meant to be cheap, affordable and thus accessible is a system proposed by Marfia et al[27]. Designed to combat work injuries – back pain and upper extremities injuries using a simple webcam, it is a demonstration of the possibilities that computer vision enabled systems have, without needing new or expensive equipment. It is useful for office jobs, where it can correct a person's posture, as well as for manual work. It has been demonstrated to correctly identify mistakes in motions that are involved in manual labor, including hammering.

Related to affordable systems like the one mentioned above is accessibility across borders, enabled by a global internet. In countries where there is a lack of trained medical professionals, healthcare could be made accessible especially using open source and free programs.[28] This software sometimes requires specialized hardware, which does not have to be expensive either. An ocular fundus camera powered by a small and cheap computer has been introduced by researchers of the University of Illinois.[29]

Computer vision also has the potential to be helpful in medical research, including the discovery of new drugs. In one critical step in the process, called protein crystallization, missing protein crystals, which are difficult to find, can result in a lost opportunity. A neural network can increase the likelihood of finding a protein.[30]

3.1.3 Mobile Devices

The proliferation of mobile devices such as smartphones in the last decade has permitted users to start doing many tasks on their internet-connected devices. Their ever-greater computational powers coupled with constantly improving camera and other

[26] (cf. Poplin et al., 2018)

[27] (cf. Marfia and Roccetti, 2017)

[28] (cf. Jones et al., 2018, p. 224)

[29] (cf. Shen and Mukai, 2017)

[30] (cf. Bruno et al., 2018)

sensors could have an impact in healthcare. Many phones now include a dedicated Health application that, amongst other things, aim to unite information about their users' health in one place.[31] This includes information that is directly measured by the phone or devices connected to it such as steps walked, flights climbed (measured by the phone's sensors or connected wearable devices), weight (collected from a connected scale), heart rate (collected from a wearable device like the Apple Watch) and even information from EHR systems deployed in hospitals. A whole new industry has been born on the premise of measuring data about users' activity with companies like Fitbit selling dedicated trackers. Computer vision could, coupled with other AI-systems interpreting this data, provide users with real-time monitoring of their health and potentially answer queries regarding their bodies by analyzing pictures taken (e.g. a dark spot on their skin).[32]

All of the examples outlined above show an important distinction that has to be made between AI that is designed for the medical field and other forms of AI. AI in medicine works together with doctors, delivering useful results and giving them a second opinion on the matter at hand.[33]

3.2 Computer Vision in the Car Industry

Self-driving vehicles will most likely be the first physical manifestation of AI that the general public will be introduced to. A consequence of that will be large portions of the population associating the term “artificial intelligence” especially with cars and other vehicles. It is therefore necessary for creators of self-driving systems to act in a responsible way not only in order not to threaten the lives of people, but also because the experience passengers gain will largely be associated with AI as a whole.[34]

Computer vision is the backbone of an autonomous car's decision-making system. While other communication channels are needed, such as vehicle-to-vehicle (V2V)

[31] (cf. Apple Inc., n.d.)

[32] (cf. Stone et al., 2016, p. 29)

[33] (cf. Chen, 2013)

[34] (cf. Stone et al., 2016, p. 19)

and vehicle-to-infrastructure (V2I) communication, it is computer vision that is most critical.[35]

3.2.1 Advantages

Automated driving is expected to improve safety[36], optimize traffic flow, reduce CO_2-emissions and fuel consumption and enhance the mobility of elderly people and unconfident drivers. Cities may become less congested as the need to own a car diminishes. As with every form of automation, the cost of "artificial drivers" is much lower than the cost of a human driver. Hence, a ride-sharing company like Uber can offer their services at a lower price. This change will make it no longer practical or necessary to own a car, freeing up parking places and roads. In turn, data gathered from these autonomous cars can be used to further optimize traffic, by considering which routes individual users frequently take and potentially grouping together similar routes for a cheaper price.[37] The privacy implications thereof are discussed in chapter 4.1.

3.2.2 Self-driving Vehicles Currently in Existence

Waymo, an Alphabet subsidiary, is one of the companies turning self-driving vehicles into reality. Their cars, semi-autonomous with human drivers acting as a fallback, have driven more than 10 million miles in total.[38]

Waymo describes the many ways it ensures safety in its publicly available safety report. Its cars are equipped with many redundant systems that are intended to overtake the default ones in the event of a failure. This includes, amongst others, Backup Computing, Backup Braking, Backup Steering and Backup Power Systems. These are measures to ensure that the self-driving car is indeed safer than human drivers are. The most important part, however, is the actual self-driving capability. This is first trained in a simulation, then tested in a closed course and lastly validated in the real world.[39]

[35] (cf. Watzenig and Horn, 2017)

[36] (cf. Weiland, 2017)

[37] (cf. "UberPool vs. UberX - How Does UberPool Work?," n.d.)

[38] (cf. Krafcik, 2018)

[39] (cf. Waymo, 2017)

In 2016, a car manufactured by the automaker Tesla and equipped with its “Autopilot” technology failed to brake while in self-driving mode. Incidents like this could have an adverse effect on the public’s perception of self-driving vehicles, demonstrating that the system is not infallible.[40] More importantly, the first self-driving car death is attributed to Uber’s fully self-driving fleet, which is trained by the company’s employees. They intervene only in edge cases, when it is absolutely needed.[41] On March 18, 2018, a 49-year-old woman was killed while crossing the street, the driver failed to react in time.[42]

A recent prediction suggests that autonomous trucks could be introduced to the road in as soon as 5 to 10 years.[43] This presents a direct threat to almost 3.3 million truck drivers’ jobs just in the United States.[44] [45]

[40] (cf. Vlasic and Boudette, 2017)

[41] (cf. Brubaker, 2018)

[42] (cf. Levin, 2018)

[43] (cf. Freedman, n.d.)

[44] (cf. “Delivery Truck Drivers and Driver/Sales Workers,” n.d.)

[45] (cf. “Heavy and Tractor-trailer Truck Drivers,” n.d.)

4 Possible Concerns About AI

4.1 Surveillance and Privacy

4.1.1 The Importance of Privacy

The award-winning journalist Glenn Greenwald, who led the *Guardian* team reporting on the NSA's overarching surveillance programs, describes the importance of privacy in his book, *No Place to Hide*. Not only does he provide anecdotal evidence by pointing out that even anti-privacy advocates put locks on their bathroom doors and do not reveal intimate information to anyone who asks for it, revealing the hypocrisy of their claims, Greenwald also explains how a lack of privacy can present a serious threat to democracies. He stresses that if an individual's every move is closely monitored, the range of actions they consider is dramatically narrower. Thus, a person living like this, by definition, cannot be free. Totalitarian regimes around the world use this to their advantage and make sure to have wide-spread surveillance in place in order to enforce absolute compliance from their citizens.[46]

The American Civil Liberties Union (ACLU) identified 4 reasons why public video surveillance is not worth the risks. Those include lack of evidence of actually reducing crime, susceptibility to abuse by criminals, public institutions, individuals etc., lack of limits or controls and changes in behavior that are caused by the awareness of constant surveillance.[47]

According to Neil M. Richards, Professor of Law at Washington University in St. Louis, surveillance is harmful in two ways. First, surveillance is harmful because it prevents people from experimenting with new and controversial ideas. Second, the disparity between the watcher and the watched can result in various harms, including discrimination and coercion.[48]

As we will see, artificial intelligence can empower both government surveillance and data collection conducted by private companies.

[46] (cf. Greenwald, 2014, p. 170)

[47] (cf. ACLU, n.d.)

[48] (cf. Richards, 2013)

4.1.2 Automated Surveillance by Governments, Facial Recognition and CCTV Cameras

In a world using AI, the existing infrastructure can be used in new ways. For example, closed-circuit television (CCTV) cameras already installed in a city combined with facial recognition algorithms pose a new kind of threat to citizens' privacy. Footage obtained by those cameras can now be analyzed using AI, therefore making it possible to build databases with detailed information on the whereabouts of ordinary people. This is not possible without AI, because it would be far too complex for people alone to achieve, as it would require continuously looking at video streams from all cameras while at the same time remembering faces seen before.

In the United States, the "Integrated Electronic Security and Surveillance System, Command Communication and Control System of Systems" (IESS/C3) was a plan to equip New York City subway stations with cameras feeding data back into an AI-powered system capable of detecting threats to New Yorkers. It was announced in 2005. When the system was supposed to be finished, only a few parts of it were ready. "What appears technologically feasible at the planning stage can apparently be undermined by the most ordinary reasons in practice." McClain cites changes in the background of an image as possible causes for failure.[49]

However, a 2018 report by *The Intercept* revealed that IBM has been granted access to the New York Police Department's public cameras. Footage obtained from these cameras was then used to develop technology that allows to search by skin color, using tags such as "White", "Black" and "Asian".[50] While developed in order to make searching footage easier, labels based on skin color threaten to reinforce stereotypes that police officers have.

In addition, technology is moving forward and delivering new technology that may enable automated surveillance. This includes new products such as hard drives, specifically designed with AI-surveillance in mind.[51] [52]

[49] (cf. McClain, 2018)

[50] (cf. Joseph and Lipp, 2018)

[51] (cf. Business Wire, 2018)

[52] (cf. Business Wire, 2017)

The Chinese government uses cameras combined with facial recognition in its Xinjiang region, creating a virtual "prison" that alerts authorities "when targeted people venture more than 300 meters (1,000 feet) beyond designated 'safe areas'", Bloomberg reported in January 2018.[53] Combined with the so-called "Social Credit System", every citizen's trustworthiness can be expressed by a number.[54]

"Trustworthiness" in this context is to be understood as a subjective rating based on the Chinese government's preferences. Being a country rated by Freedom House as "Not Free"[55] as well having a record of various different human rights violations as reported in the 2017 Country Reports on Human Rights Practices released by the U.S. Department of State[56] indicates that the Chinese government will use this score for nefarious purposes such as suppressing dissent and securing its power.

The European Union has funded a project named iBorderCtrl (Intelligent Portable Control System). Project participants include the Hungarian National Police and the State Border Guard of the Republic of Latvia. Its goal is to enable faster border control. Part of iBorderCtrl is the Automatic Deception Detection System (ADDS) which "quantifies the probability of deceit in interviews by analysing interviewees non-verbal micro expressions".[57]

4.1.3 Data Collection by Private Companies

Several AI-enabled services are in use today, available to consumers. Companies such as Apple[58], Samsung[59], Google[60], Amazon[61] and Microsoft[62] all offer "smart assistant" technology. The idea is that a dedicated speaker or software program on a

[53] (Bloomberg News, 2018)

[54] (cf. Botsman, 2017)

[55] (cf. Freedom House, 2018)

[56] (cf. US Department of State, Bureau of Democracy, Human Rights and Labor, 2017)

[57] ("Technical Framework | iBorderCtrl," n.d.)

[58] (cf. "Siri," n.d.)

[59] (cf. "Bixby," n.d.)

[60] (cf. "Google Assistant - Just Say 'Hey Google' and Make Google Do It," n.d.)

[61] (cf. "Echo & Alexa - Amazon Devices - Amazon Official Site," n.d.)

[62] (cf. "What is Cortana?," n.d.)

device (e.g. a phone) is listening for a keyword to trigger its function. These are “Hey Siri”, “Alexa” and “Hey Google” or “OK, Google” for Apple’s, Amazon’s and Google’s assistants respectively. A user can then use their assistant to request information such as weather forecasts, sports results or personalized requests like received messages. They can also prompt the assistant to order pizza, hail a cab or ask it to do many other things.[63]

While these assistants may be convenient to a certain extent, they also tend to collect large amounts of data on usage and their users’ preferences. Google, for instance, is financially dependent on ad revenue generated from personalized ads, creating an incentive to collect as much data as possible.[64] Meanwhile, consumers rarely focus on features like privacy when buying a new product or are not even aware of the information that is collected. For instance, iRobot, a company selling “autonomous” vacuum cleaners, was creating maps of users’ homes that, additionally to improving the vacuum cleaner’s function, can be sold to advertisers.[65] While this was allowed in their privacy policy, the complicated texts written using legal jargon can be incomprehensible to the average user. GDPR, an EU regulation described below, aims to fix this by requiring companies to inform users about the specific uses of their data in plain language.[66] Users can also delete their data at alexa.amazon.com, view and delete earlier Google Assistant requests at myactivity.google.com. For Apple’s Siri, turning the assistant off in a device’s Settings app will delete some data. According to a privacy statement found under the “Siri & Search” section in an iOS device’s Settings app running version 12.1.1, this will delete User Data. Some information “that has been disassociated from [a user] may be retained”.

However, even in cases where consumers trust corporations with their data, security vulnerabilities can give third parties information in the form of metadata. The term metadata is generally understood as information about information, for example the time and duration of a phone call as opposed to the content of it.

[63] (cf. “Google Assistant - What can your Assistant do?,” n.d.)

[64] (cf. Rosenberg, 2015)

[65] (cf. Jones, 2017)

[66] (cf. GDPR, 2016)

Amazon envisions their Echo suite of products to be the basis of a smart home, i.e. the medium through which users communicate with the home. Adversaries or anyone who has access to a home's encrypted traffic, including Internet service providers (ISPs), could then gain information about users and would not even have to circumvent encryption in transit. It is enough that they can see metadata about when a user interacted with a device, because this generates sudden spikes in traffic. Examples can be found in the image below. Section A shows traffic spikes generated by a sleep monitor. We can recognize 3 major spikes that are tied to 3 events, namely going to bed, getting out of bed during the night and getting up in the morning. Sections B and C demonstrate how an ISP might be able to tell whether or not someone is at the user's home. Since the Nest Camera only transmits a video feed when motion is detected, low traffic means no motion is detected. Sections D and E describe how traffic reveals that a device such as a WeMo Switch and an Amazon Echo have been directly interacted with, in the case of the Echo this means posing a question to the smart assistant. Thus, metadata can provide an observer with detailed information, including when a user gets up in the morning and leaves for work.

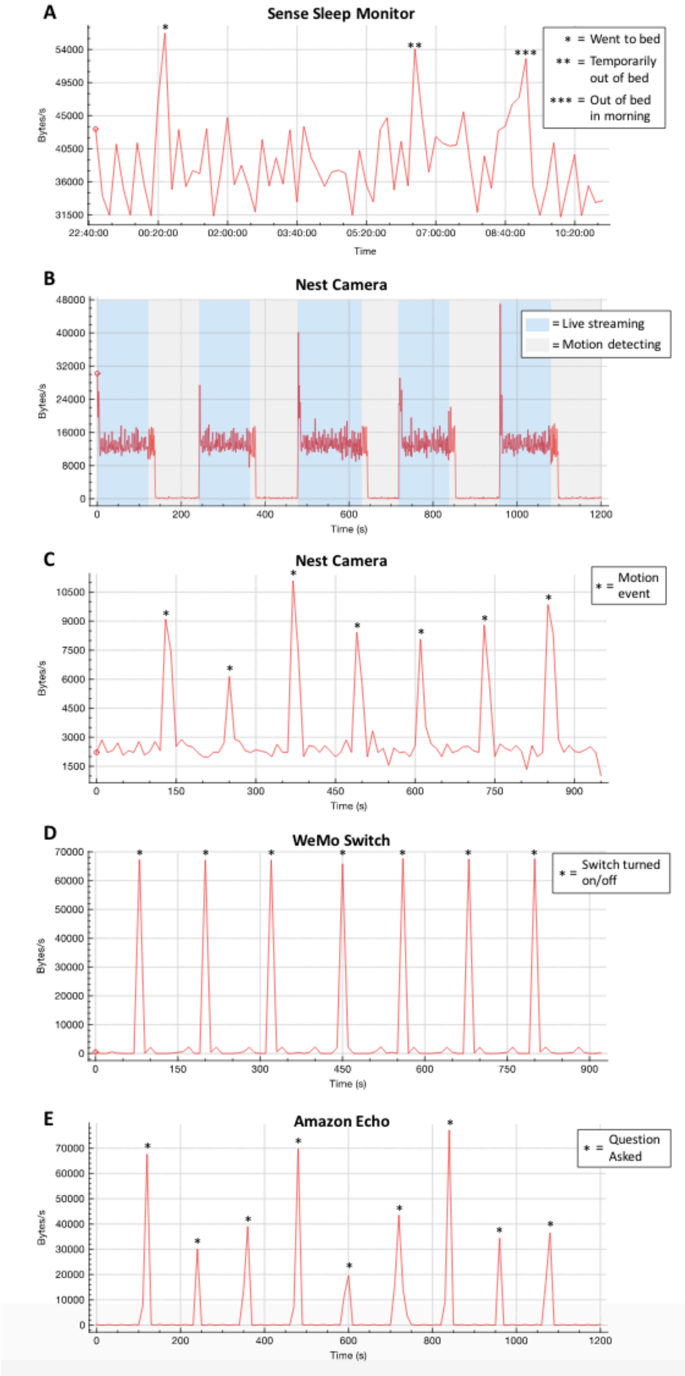

Figure 1 (Apthorpe et al.)

Moreover, this is not limited to a particular company's products but is a natural consequence of the way that various internet protocols work. Apthorpe et al. demonstrated this to be true by analyzing smart devices' internet traffic without employing deep packet inspection, i.e. analyzing the encrypted data itself. It was possible to tell smart devices apart based on IP addresses they communicated with and DNS queries they made.[67] The article then goes on to explain what can be done in order to improve users' privacy. This includes regulatory rules limiting ISPs' data collection. GDPR is helpful in this case, requiring companies to gain explicit consent from users before gathering data on them. The article also points out that data collected from devices related to security and healthcare is potentially more sensitive than devices like the Echo, giving the example of a sleep monitor, where internet traffic is related to when a user sleeps. An ISP could infer from this data that the user suffers from sleeping disorders, information that is sensitive and could be attractive to third parties seeking to sell services or products related to improving their customers' sleep.

4.1.4 Europe

In Europe, the EU General Data Protection Regulation, or GDPR, is a regulation enforced across all EU Member States.

> *"The new law covers the personal data of all EU residents, regardless of the location of the processing. Personal data is information that, directly or indirectly, can identify an individual, and specifically includes online identifiers such as IP addresses, cookies and digital fingerprinting, and location data that could identify individuals. This is much wider than the concept of personally identifiable information under US privacy law."* [68]

The European Union thus now expects software developers to implement privacy protections by default. This is most pronounced in a new key concept called Data Minimization, meaning having to use the absolute minimum of data required to complete a task[69], which further reduces the amount of information that potential AI developers will have access to. This, on one hand, improves consumer privacy, but on the other hand, it could potentially present a burden too difficult for European companies developing AI to overcome. That would in turn make AI software developed in the United States and in China, where regulation is less strict, more advanced and prevalent

[67] (cf. Apthorpe et al., 2017, p. 2)

[68] (Goddard, 2017)

[69] (cf. GDPR, 2016, Article 5)

worldwide. GDPR also requires companies to provide users with options such as requesting a copy of the data that the company keeps on them and requesting this data to be deleted. It is thus a way for users to "opt out" of any surveillance conducted by private companies at any time. This, however, does not apply to government surveillance.

In France, a parliamentary mission assigned by the Prime Minister Édouard Philippe produced the "AI for Humanity" report, which was written as a guidance to the French president in his policies regarding AI. On March 29, 2018, he "presented his vision and strategy to make France a leader in artificial intelligence".[70] The report sets out ways to create AI that is beneficial to the human population, what it calls "meaningful AI". Amongst other things, "meaningful" refers to taking the approach of "opening up the black box", i.e. creating AI in a way that is transparent.

It shows how European-developed AI can preserve its users' privacy by complying to the GDPR while simultaneously not lagging behind "American" or "Chinese" AI in its capabilities.[71] This includes the creation of a "European Data Ecosystem", in which data would be considered a shared common resource available to everyone.

4.1.5 United States

As of December 22, 2018, a federal law to protect consumers' privacy does not yet exist in the United States.[72] Some states have their own privacy laws, California passed one in June 2018, called the California Consumer Privacy Act of 2018. The legislation has been widely compared to Europe's GDPR.

> *"Companies that store large amounts of personal information — including major players like Google and Facebook — will be required to disclose the types of data they collect, as well as allow consumers to opt out of having their data sold."*[73]

[70] ("AI for humanity," n.d.)

[71] (cf. Villani, 2018b)

[72] (cf. Kelly, 2018)

[73] (Lecher, 2018, para. 2)

The revelation in early 2018 that a British political consulting firm called Cambridge Analytica used information gathered from Facebook for political purposes fueled debates about a possible introduction of a federal privacy law in the US.[74]

> *"'If Facebook and other online companies will not or cannot fix these privacy invasions, then we will,' warned Sen. Bill Nelson (Fla.), the ranking Democrat on the Commerce Committee."* [75]

Parallels can be found between the Cambridge Analytica scandal and new legislation proposed. The CONSENT act proposed in April 2018 would require explicit permission from users for the use of sensitive information.[76]

Contrary to GDPR, the scope of what is considered to be sensitive information is much narrower. Where GDPR says anything that can identify an individual is sensitive, the CONSENT act's definition only encompasses information such as Social Security numbers, health and financial information.[77] The Social Media Privacy Protection and Consumer Rights Act of 2018 extends the definition to e-mail addresses and telephone numbers.[78]

These new weaker privacy protections could override stronger state-level protections like the California Consumer Privacy Act of 2018 if they went came into force.[79]

Stanford's AI100 report recommends to "remove the perceived and actual impediments to research on the fairness, security, privacy, and social impacts of AI systems".[80] It suggests that AI systems should be reverse-engineerable, similar to the French AI for Humanity recommendation of creating transparent AI.

[74] (cf. Cadwalladr and Graham-Harrison, 2018)

[75] (Sullivan, 2018, para. 9)

[76] (cf. Inside Privacy, 2018)

[77] (cf. Markey, n.d., p. 4)

[78] (cf. Kennedy and Klobuchar, 2018)

[79] (cf. Guliani, n.d.)

[80] (Stone et al., 2016, p. 43)

4.1.6 Conclusion

Ethical issues regarding privacy protections are one of the most urgent ones to be tackled. Artificial intelligence enables automated surveillance by governments, examples of which can be found in the United States, China and Europe, as described in more detail above. Similarly, there is growing concern regarding data collection by private companies. As homes become more proliferated with smart devices, regulation that protects consumers' right to data protection and privacy are needed. This is in part because encryption itself cannot prevent all kinds of surveillance, since metadata can reveal much information too. In Europe, this kind of regulation comes in the form of GDPR, while similar but generally weaker bills have been proposed in the United States.

4.2 Economic Impact

Artificial intelligence is largely believed to replace human workers in different areas. Traditionally, those include types of work where physical strength is required, e.g. a factory. Highest skilled workers like doctors and researchers are initially expected to work together with AI (see chapter 3.1). There is an AI-based drug discovery program based on IBM's Watson available today.[81] This program still requires human intervention and only acts as a tool for researchers. As all machine learning based systems available today only focus on very specific tasks (narrow AI), there is a need for human workers to bring different information together and put it into the right context. A radiologist, for instance, may also have to consult with other physicians, treat diseases and tailor particular details to a patient's specific situation.[82] Still, as artificial intelligence becomes more sophisticated, doctors could eventually even be replaced.[83]

So far, technology has mainly taken jobs of middle skilled workers, e.g. travel agents with new websites that may or may not also incorporate AI. However, a new kind of

[81] (cf. “IBM Watson for Drug Discovery - Overview - United States,” n.d.)

[82] (cf. Davenport and Dreyer, 2018)

[83] (cf. Purdy and Daugherty, 2016)

jobs may be born out of necessity as a consequence of a gradual proliferation of intelligent autonomous systems, e.g. programmers, computer science experts, or people with degrees in mathematics.

The shift from human workers to computerized systems will most likely be a gradual one. As artificial intelligence takes these jobs, new economic and political measures will be taken to protect those whose jobs are endangered. [84]

"As children in traditional societies support their aging parents, perhaps our artificially intelligent "children" should support us, the "parents" of their intelligence."[85]

Universal basic income (UBI) is a social security scheme in which all citizens regularly receive a certain amount of money. They may then choose to work and make even more money or to live from the basic income. Job automation may enable this system, by making labor very cheap. UBI can also work as one of the new social safety nets for people who lost their job to a robot.

There has been a referendum on UBI in Switzerland in 2016. 77 percent of Swiss voters reject the idea.[86]

A similar proposition intended to shield workers from drastic changes is a "robot tax". This concept has been endorsed by Bill Gates[87] but rejected by the European parliament. On one hand, it may reduce inequality, but on the other hand, hinder technological progress in the EU, the latter one being the reason for its rejection.[88]

4.3 Ethical Concerns

Ethics is "the discipline concerned with what is morally good and bad, right and wrong".[89] Today, core moral principles are shared by people. They are the basic fabric that makes our societies function. When machines start doing people's jobs, they will need to be taught these principles. Determining what exactly these principles are is a

[84] (cf. Stone et al., 2016)

[85] (Stone et al., 2016, p. 39)

[86] (cf. BBC, 2016)

[87] (cf. Shiller, 2017)

[88] (cf. Kharpal, 2017)

[89] (Hosch, 2009c)

challenge in itself. Like all technology, artificial intelligence is nothing more than a tool and the impact it will have on the world is highly dependent on the people that aim to make use of it. Establishing moral principles and legislation that ought to be adhered to is therefore necessary.

4.3.1 Self-driving Vehicles

Before artificial intelligence-powered cars can hit the road, moral principles have to be established to determine such a car's "priority list" in the case of an occurrence of "the modern trolley problem":

> *"An autonomous vehicle has a brake failure, leading to an accident with inevitably tragic consequences; due to the vehicle's superior perception and computation capabilities, it can make an informed decision. Should it stay its course and hit a wall, killing its three passengers, one of whom is a young girl? Or swerve and kill a male athlete and his dog, who are crossing the street on a red light?"*[90]

It is necessary to implement a decision-making process that roughly mirrors society's preferences. Noothigattu et al. propose a system consisting of 4 steps: data collection, i.e. collecting information on what society agrees upon regarding these ethical choices, learning, creating a model that is able to predict the preferences of each voter, summarization, or combining these models into a single one which reflects choices society would make as a whole, and aggregation, running the algorithm by determining each voter's preferences and then simulating "voting", which will inform an AI system's decision.[91]

Notably, this moral principle "priority list" could backfire in the event that a self-driving car saw what it believed were children on the road, which would actually be boxes flying around. Deciding that the lives of the children are higher up on the "priority list" and swerving whereby potentially injuring or killing its occupants and/or other motorists is an event that could have been avoided if a human driver was operating the vehicle.[92]

[90] (Noothigattu et al., 2017)

[91] (cf. Noothigattu et al., 2017)

[92] (cf. Brubaker, 2018, p. 6)

A website set up by MIT has surveyed people around the world on their moral choices.[93] As the results of this "Moral Machine experiment" show, people from different cultural backgrounds can be divided into three sections based on their preferences – Awad et al. named them "Eastern", "Western" and "Southern" based on their geographical origin. Their answers can sometimes vary significantly. For example, people belonging to the "Southern" group demonstrated much greater willingness to spare the young and females than people from "Eastern" countries.[94]

As the paper points out, the ethical rules proposed in 2017 by the German Ethics Commission on Automated and Connected Driving partly align with the experiment's findings, e.g. according to German Ethical Rule number 7, the protection of human life is more important than the protection of animal life. Nonetheless, at other times the findings did not line up with the German Ethical Rules, notably with rule number 9 which prohibits features such as age from affecting the person's "priority score".

The creation of a "priority list" would also have to ensure that bigotry and biased views held by the general public is not introduced into this list.

4.3.2 Bias

Artificial intelligence has been shown to include biases[95] [96] [97] and it is therefore important to implement AI in a way that is responsible and aware of the fact that biases can be introduced by accident. This relates to the aforementioned AI for Humanity report and its suggestion of "opening up the black box". It must act so that its decision can be traced back and rationalized if needed, e.g. by law enforcement. The report also suggests that ethics should become part of a standard curriculum for future developers, so that they can incorporate it by design. The question is raised whether not to keep human judgment in certain areas, where decisions could have the most serious impact. The establishment of an ethics committee is recommended. It would be

[93] (cf. "Moral Machine," n.d.)

[94] (cf. Awad et al., 2018)

[95] (cf. Ferryman and Pitcan, 2018)

[96] (cf. Knight, 2017)

[97] (cf. Caliskan et al., 2017)

desirable to ensure that everyone is represented by this committee, creating a combination of people from different cultural and socio-economic backgrounds.

Weng et al. created 4 different machine learning algorithms that predicted cardiovascular risk in patients. These algorithms also considered factors such as ethnicity and socioeconomic status. While these factors may be useful to predict risk factors more precisely, the danger of creating a "black box", where it would no longer be clear why a certain algorithm made a particular choice, is present.[98]

In the United States, *The Guardian* reported in 2017 that a House oversight committee hearing revealed that a facial recognition database used by the FBI contains images of US adults without consent, along with an algorithm that is wrong 15% of the time and more likely to misidentify people with dark skin color.

> *"'No federal law controls this technology, no court decision limits it. This technology is not under control,' said Alvaro Bedoya, executive director of the center on privacy and technology at Georgetown Law."*[99]

Congressman Elijah Cummings expressed concern regarding the higher inaccuracy of this facial recognition system when scanning faces of African Americans. If Artificial intelligence systems such as those employed at border crossings are inherently flawed in this way, meaning they are showing bias against ethnic groups, AI could become a threat to democracy itself.

> *"As a society, we are now at a crucial juncture in determining how to deploy AI-based technologies in ways that promote, not hinder, democratic values such as freedom, equality, and transparency."* [100]

[98] (cf. Weng et al., 2017)

[99] (Solon, 2017)

[100] (Stone et al., 2016)

4.3.3 Accountability and Liability

As artificial intelligence transforms many industries, responsibilities will shift accordingly. In AI-enabled self-driving vehicles, the question of responsibility becomes very urgent. In the case of an accident, manufacturers of self-driving systems will be held accountable when the car was at fault.[101] A similar situation arises when a patient is misdiagnosed by artificial intelligence. It is not clear if the company that created the system or the hospital where it was used is at fault.

According to case law, doctors are allowed to make mistakes when their decisions are made in good faith. It would be reasonable to allow for a similar margin of error when these same decisions are made by artificial intelligence. AI systems are trained on a model provided by people. Errors will unavoidably occur.

[101] (cf. Oberly, 2017)

5 Conclusion

Artificial intelligence is the capability of a computer to act intelligently. Computer vision, a subfield of artificial intelligence, is concerned with replicating the human visual system. Most relevant is object detection, where a computer recognizes an object in an image, locates and names it.

In medicine, highly specialized computer vision models are trained on existing data which has been recorded in an electronic health record system. These models can then be used directly by doctors or distributed over the Internet. Together with cheap hardware, they can make medical diagnosis accessible in new parts of the world. Mobile devices running similar software monitor a user’s health in real time.

Computer vision facilitates a self-driving car’s understanding of the world. In the United States, companies like Waymo and Uber are already testing them, some on public roads. The goal of these efforts is to create a better driver to reduce death on roads.

Artificial intelligence is nothing more than a tool. Therefore, it can be equally used for nefarious purposes. Governments and companies motivated by financial gain employ AI for surveillance purposes. In the case of surveillance conducted by private companies, regulation is necessary to protect consumers. GDPR fulfills this function in Europe. Similar federal legislation has been proposed in the United States. When it comes to government surveillance, the same laws do not apply.

AI will replace people’s jobs, initially physical work, but later on “mental” work done by highly trained professionals could be replaced as well. Proposals to combat this include universal basic income and the taxation of robots.

At last, these new “worker robots” need to know people’s ethical principles. Ethical principles can vary from person to person, geographical origin also plays a role. Biases learnt from people must not be part of a robot’s understanding of the world. In the case of a failure, the company who created the AI will be sometimes held accountable. However, this will not always be the case, as a margin of error is to be expected in every decision-making system.

6 Bibliography

6.1 Primary Literature

Adams, J., Khan, H.T.A., Raeside, R., 2014. Research Methods for Business and Social Science Students, Second edition. ed. SAGE Publications Pvt. Ltd, Los Angeles.

Chen, C.H., 2013. Computer Vision In Medical Imaging. World Scientific, Hackensack, NJ.

Greenwald, G., 2014. No Place to Hide: Edward Snowden, the NSA, and the U.S. Surveillance State. Metropolitan Books, New York, NY.

Watzenig, D., Horn, M., 2017. Automated Driving: Safer and More Efficient Future Driving. Springer International Publishing Imprint: Springer, Cham.

Yoshida, S.R., 2011. Computer Vision, Computer Science, Technology and Applications. Nova Science Publishers, Inc, Hauppauge, N.Y.

6.2 Journal Articles

Apthorpe, N., Reisman, D., Feamster, N., 2017. A Smart Home is No Castle: Privacy Vulnerabilities of Encrypted IoT Traffic. ArXiv170506805 Cs.

Awad, E., Dsouza, S., Kim, R., Schulz, J., Henrich, J., Shariff, A., Bonnefon, J.-F., Rahwan, I., 2018. The Moral Machine experiment. Nature 563, 59–64. https://doi.org/10.1038/s41586-018-0637-6

Bazzoli, F., 2018. Cloud-based platform enables use of AI on medical images. Healthdatamanagement.com.

Bruno, A.E., Charbonneau, P., Newman, J., Snell, E.H., So, D.R., Vanhoucke, V., Watkins, C.J., Williams, S., Wilson, J., 2018. Classification of crystallization outcomes using deep convolutional neural networks. PLOS ONE 13, e0198883. https://doi.org/10.1371/journal.pone.0198883

Caliskan, A., Bryson, J.J., Narayanan, A., 2017. Semantics derived automatically from language corpora contain human-like biases. Science 356, 183–186. https://doi.org/10.1126/science.aal4230

Davenport, T.H., Dreyer, K.J., 2018. AI Will Change Radiology, but It Won’t Replace Radiologists. Harv. Bus. Rev. Digit. Artic. 2–5.

Felzenszwalb, P.F., Girshick, R.B., McAllester, D., Ramanan, D., 2010. Object Detection with Discriminatively Trained Part-Based Models. IEEE Trans. Pattern Anal. Mach. Intell.

32, 1627–1645. https://doi.org/10.1109/TPAMI.2009.167

Ferryman, K., Pitcan, M., 2018. Fairness in Precision Medicine 54.

Girshick, R., Donahue, J., Darrell, T., Malik, J., 2013. Rich feature hierarchies for accurate object detection and semantic segmentation. ArXiv13112524 Cs.

Goddard, M., 2017. The EU General Data Protection Regulation (GDPR): European regulation that has a global impact. Int. J. Mark. Res. 59, 703–705. https://doi.org/10.2501/IJMR-2017-050

Huang, T.S., 1996. Computer Vision: Evolution and Promise 5.

Jones, L.D., Golan, D., Hanna, S.A., Ramachandran, M., 2018. Artificial intelligence, machine learning and the evolution of healthcare: A bright future or cause for concern? Bone Jt. Res. 7, 223–225. https://doi.org/10.1302/2046-3758.73.BJR-2017-0147.R1

Jones, M.J., Viola, P., n.d. Robust Real-time Object Detection 30.

Marfia, G., Roccetti, M., 2017. A practical computer based vision system for posture and movement sensing in occupational medicine. Multimed. Tools Appl. 76, 8109–8129. https://doi.org/10.1007/s11042-016-3469-0

McClain, N., 2018. The horizons of technological control: automated surveillance in the New York subway. Inf. Commun. Soc. 21, 46–62. https://doi.org/10.1080/1369118X.2016.1260624

Noothigattu, R., Gaikwad, S. “Neil” S., Awad, E., Dsouza, S., Rahwan, I., Ravikumar, P., Procaccia, A.D., 2017. A Voting-Based System for Ethical Decision Making. ArXiv170906692 Cs.

Oberly, D.J., 2017. Changing Lanes: Determining Autonomous Car Liability: Who is responsible when self-driving cars are in an accident? Claims 65, 39–41.

Paul-Choudhury, S., 2016. Outsmarted? New Sci. 230, 18–19. https://doi.org/10.1016/S0262-4079(16)31133-2

Poplin, R., Varadarajan, A.V., Blumer, K., Liu, Y., McConnell, M.V., Corrado, G.S., Peng, L., Webster, D.R., 2018. Prediction of cardiovascular risk factors from retinal fundus photographs via deep learning. Nat. Biomed. Eng. 2, 158–164. https://doi.org/10.1038/s41551-018-0195-0

Richards, N.M., 2013. The Dangers of Surveillance. Harvard Law Review.

Sánchez, J., Perronnin, F., Mensink, T., Verbeek, J., 2013. Image Classification with the Fisher Vector: Theory and Practice. Int. J. Comput. Vis. N. Y. 105, 222–245. http://dx.doi.org.ezproxy.techlib.cz/10.1007/s11263-013-0636-x

Shen, B.Y., Mukai, S., 2017. A Portable, Inexpensive, Nonmydriatic Fundus Camera Based on the Raspberry Pi® Computer. Journal of Ophthalmology. https://doi.org/10.1155/2017/4526243

Weng, S.F., Reps, J., Kai, J., Garibaldi, J.M., Qureshi, N., 2017. Can machine-learning improve cardiovascular risk prediction using routine clinical data? PLOS ONE 12, e0174944. https://doi.org/10.1371/journal.pone.0174944

6.3 Laws, Bills and Regulations

GDPR, 2016. Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation) (Text with EEA relevance). OJ L.

Kennedy, J., Klobuchar, A., 2018. Social Media Privacy Protection and Consumer Rights Act of 2018. URL https://www.govtrack.us/congress/bills/115/s2728/text (accessed 2.6.19)

Markey, n.d. Customer Online Notification for Stopping Edge-provider Network Transgressions. URL https://www.markey.senate.gov/imo/media/doc/CONSENT%20Act%20text.pdf (accessed 2.6.19)

6.4 Reports

Hundred Year Study on Artificial Intelligence: Report of the 2015-2016 Study Panel. Stanford University, Stanford, CA.

Purdy, M., Daugherty, P., 2016. Why Artificial Intelligence Is The Future Of Growth. Accenture.

Stone, P., Brooks, R., Brynjolfsson, E., 2016. “Artificial Intelligence and Life in 2030.” One

US Department of State, Bureau of Democracy, Human Rights and Labor, 2017. CHINA (INCLUDES TIBET, HONG KONG, AND MACAU) 2017 HUMAN RIGHTS REPORT.

Villani, C., 2018a. What Is Artificial Intelligence? Villani Mission on Artificial Intelligence.

Villani, C., 2018b. For a Meaningful Artificial Intelligence: Towards a French and European Strategy.

Waymo, 2017. Waymo Safety Report: On the Road to Fully Self-Driving.

6.5 Conference Papers

Redmon, J., Divvala, S., Girshick, R., Farhadi, A., 2016. You Only Look Once: Unified, Real-

Time Object Detection, in: 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR). Presented at the 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), IEEE, Las Vegas, NV, USA, pp. 779–788. https://doi.org/10.1109/CVPR.2016.91

6.6 Theses

Brubaker, K., 2018. Artificial Intelligence: Issues of Consumer Privacy, Industry Risks, and Ethical Concerns (M.S.). Utica College, United States -- New York.

6.7 Lexica

Hosch, W.L., 2009a. artificial intelligence (AI). Encyclopædia Britannica. *Britannica Academic*. URL https://academic-eb-com.ezproxy.techlib.cz/levels/collegiate/article/artificial-intelligence/9711 (accessed 2.6.19).

Hosch, W.L., 2009b. machine learning. Encyclopædia Britannica. *Britannica Academic*. URL https://academic-eb-com.ezproxy.techlib.cz/levels/collegiate/article/machine-learning/474180 (accessed 2.6.19)

Hosch, W.L., 2009c. ethics. Encyclopædia Britannica. *Britannica Academic*. URL https://academic-eb-com.ezproxy.techlib.cz/levels/collegiate/article/ethics/106054 (accessed 2.6.19)

6.8 Newspaper and Magazine Articles

BBC, 2016. Swiss voters reject basic income plan. BBC.

Bloomberg News, 2018. China Uses Facial Recognition to Fence In Villagers. Bloomberg.

Botsman, R., 2017. Big data meets Big Brother as China moves to rate its citizens. Wired UK.

Cadwalladr, C., Graham-Harrison, E., 2018. Revealed: 50 million Facebook profiles harvested for Cambridge Analytica in major data breach. The Guardian.

Campaign Asia-Pacific, 2015. Sir TIM BERNERS-LEE. Campaign Asia-Pacific 60–63.

Guliani, N.S., n.d. The tech industry is suddenly pushing for federal privacy legislation. Watch out. Washington Post.

Jones, R., 2017. Roomba’s Next Big Step Is Selling Maps of Your Home to the Highest Bidder. Gizmodo.

Joseph, G., Lipp, K., 2018. IBM Used NYPD Surveillance Footage to Develop Technology That Lets Police Search by Skin Color. The Intercept.

Kahn, J., 2002. It's Alive! WIRED.

Kelly, M., 2018. How Congress could rein in Google and Facebook. The Verge.

Kharpal, A., 2017. Bill Gates wants to tax robots, but the EU says, "no way, no way." CNBC.

Knight, W., 2017. Google's AI chief says forget Elon Musk's killer robots, and worry about bias in AI systems instead. MIT Technology Review.

Lecher, C., 2018. California just passed one of the toughest data privacy laws in the country. The Verge.

Levin, S., 2018. Video released of Uber self-driving crash that killed woman in Arizona. The Guardian.

Quinn, B., 2016. Google given access to healthcare data of up to 1.6 million patients. The Guardian.

Shiller, R., 2017. Why robots should be taxed if they take people's jobs. The Guardian.

Solon, O., 2017. Facial recognition database used by FBI is out of control, House committee hears. The Guardian.

Sullivan, M., 2018. Members of Congress can't possibly regulate Facebook. They don't understand it. Washington Post.

Vlasic, B., Boudette, N.E., 2017. Self-Driving Tesla Was Involved in Fatal Crash, U.S. Says. N. Y. Times.

Weiland, J., 2017. How Safe Are Self-Driving Cars? Huffington Post.

6.9 Online Sources

About Us [WWW Document], n.d. DeepMind. URL https://deepmind.com/about/ (accessed 9.2.18).

AI for humanity [WWW Document], n.d. URL https://www.aiforhumanity.fr (accessed 11.3.18).

Apple Inc., n.d. iOS - Health [WWW Document]. Apple. URL https://www.apple.com/ios/health/ (accessed 1.2.19).

Bixby [WWW Document], n.d. Samsung Electron. Am. URL /us/explore/bixby/ (accessed 11.21.18).

Delivery Truck Drivers and Driver/Sales Workers : Occupational Outlook Handbook: : U.S. Bureau of Labor Statistics [WWW Document], n.d. URL https://www.bls.gov/ooh/transportation-and-material-moving/delivery-truck-drivers-and-driver-sales-workers.htm (accessed 9.4.18).

Echo & Alexa - Amazon Devices - Amazon Official Site [WWW Document], n.d. URL https://www.amazon.com/Amazon-Echo-And-Alexa-Devices/b?ie=UTF8&node=9818047011 (accessed 11.21.18).

Freedman, D.H., n.d. If automation is already messing with our economy and our politics, just wait until self-driving trucks arrive [WWW Document]. MIT Technol. Rev. URL https://www.technologyreview.com/s/603493/10-breakthrough-technologies-2017-self-driving-trucks/ (accessed 9.4.18).

Freedom House, 2018. China | Freedom in the World 2018 [WWW Document]. URL https://www.freedomhouse.org/sites/default/files/inline_images/Metatag_IMG_Freedom_in_the_world2018_Scores_China.png (accessed 12.22.18).

Google Assistant - Just Say "Hey Google" Make Google Do It. [WWW Document], n.d. URL https://assistant.google.com/ (accessed 11.21.18).

Google Assistant - What can your Assistant do? [WWW Document], n.d. URL https://assistant.google.com/explore (accessed 11.21.18).

Heavy and Tractor-trailer Truck Drivers : Occupational Outlook Handbook: : U.S. Bureau of Labor Statistics [WWW Document], n.d. URL https://www.bls.gov/ooh/transportation-and-material-moving/heavy-and-tractor-trailer-truck-drivers.htm (accessed 9.4.18).

IBM Watson for Drug Discovery - Overview - United States [WWW Document], n.d. URL https://www.ibm.com/us-en/marketplace/ibm-watson-for-drug-discovery#product-header-top (accessed 6.11.18).

Moral Machine [WWW Document], n.d. Moral Mach. URL http://moralmachine.mit.edu (accessed 8.27.18).

Rosenberg, E., 2015. How Google Makes Money (GOOG) [WWW Document]. Investopedia. URL https://www.investopedia.com/articles/investing/020515/business-google.asp (accessed 11.21.18).

Siri [WWW Document], n.d. Apple. URL https://www.apple.com/siri/ (accessed 11.21.18).

Technical Framework | iBorderCtrl [WWW Document], n.d. URL https://www.iborderctrl.eu/Technical-Framework (accessed 12.22.18).

UberPool vs. UberX - How Does UberPool Work? | Uber [WWW Document], n.d. URL https://www.uber.com/en-CZ/ride/uberpool/ (accessed 9.4.18).

What is Cortana? [WWW Document], n.d. URL https://support.microsoft.com/en-us/help/17214/windows-10-what-is (accessed 11.21.18).

6.10 Blog Posts

ACLU, n.d. What's Wrong With Public Video Surveillance? American Civil Liberties Union. URL https://www.aclu.org/other/whats-wrong-public-video-surveillance (accessed 6.11.18).

Business Wire, 2017. Seagate Launches First Drive for AI-Enabled Surveillance. Business Wire. URL https://www.businesswire.com/news/home/20171028005001/en/ (accessed 8.21.18).

Business Wire, 2018. Western Digital Enables Artificial-Intelligence-Powered Video Surveillance with New High-Capacity Products. Business Wire. URL https://www.businesswire.com/news/home/20180619005224/en/ (accessed 8.21.18).

Inside Privacy, 2018. Senate Democrats Propose CONSENT Act. Inside Privacy. URL https://www.insideprivacy.com/united-states/congress/senate-democrats-propose-consent-act/ (accessed 11.3.18).

Krafcik, J., 2018. Where the next 10 million miles will take us. Waymo. URL https://medium.com/waymo/where-the-next-10-million-miles-will-take-us-de51bebb67d3 (accessed 2.9.19)

6.11 Images

Apthorpe, N., Reisman, D., Feamster, N., 2017. A Smart Home is No Castle: Privacy Vulnerabilities of Encrypted IoT Traffic. ArXiv170506805 Cs.

7 Table of Figures

Eidesstattliche Erklärung

Ich erkläre, dass ich die vorwissenschaftliche Arbeit eigenständig angefertigt und nur die im Literaturverzeichnis angeführten Quellen und Hilfsmittel benutzt habe.
Prag, am 13. Februar 2019

____
Unterschrift