



**Artificial Intelligence and Data Protection:
Delivering Sustainable AI Accountability in Practice**

First Report:
**Artificial Intelligence and
Data Protection in Tension**

October 10, 2018

Dear Colleagues:

I am delighted to provide you with this first report from the Centre for Information Policy Leadership's project on **Artificial Intelligence and Data Protection**.

In this report, we attempt to describe in clear, understandable terms: (1) what AI is and how it is being used all around us today; (2) the role that personal data plays in the development, deployment and oversight of AI; and (3) the opportunities and challenges presented by AI to data protection laws and norms.

We intend for this report to provide a level-setting backdrop for the next phase of our project—namely, working with data protection officials, industry leaders and others to identify practical ways of addressing challenges and harnessing the opportunities presented by AI and data protection. Our research to date suggests that those will include identifying best practices that organisations are already employing to ensure not only legal compliance, but also legal and ethical accountability when using personal data with AI. And they will include important ways of interpreting and applying existing data protection laws to protect privacy without unnecessarily stifling adoption of and innovation in AI, as well as considerations for future data protection laws.

We are grateful to CIPL members as well as academic and government experts for their participation in this first document, and we look forward eagerly to collaborating further on the critical effort to identify solutions and practical tools in our forthcoming report.

CIPL is often described as a bridge among diverse constituencies in the pursuit of rational, accountable, effective data protection and the responsible use of data. Never has that been needed more than in the context of AI, which already delivers extraordinary benefits to individuals and society, but precisely because of its power and impact requires even more collaborative efforts to ensure that it is developed and used in ways that respect personal privacy.

At CIPL, we are committed to that task. We eagerly welcome your ideas, your insights, and your partnership in our joint journey towards achieving that goal. For more information visit <https://www.informationpolicycentre.com/> or reach out to me at bellamy@huntonAK.com.

Sincerely,

Bojana Bellamy
President

I. Executive Summary

Artificial intelligence (AI) has rapidly developed in recent years. Today, AI tools are used widely by both private and public sector organisations around the globe, and governments around the world have expressed a commitment to AI's continued development. The capabilities of AI now and in the near future create widespread and substantial benefits for individuals, institutions and society.

However, these same technological innovations raise important issues, including questions about how to deliver practical compliance with data protection laws and norms when building and implementing AI technology and on the tension between AI and existing data protection legal requirements. As a result, we have both an opportunity and an obligation to develop principles, best practices and other accountability tools to encourage responsible data management practices, respect and even bolster data protection, and remove unnecessary roadblocks for the future development of these innovative technologies. Clarifying the application of existing data protection law on AI will be essential to ensuring that limited resources are not wasted on protecting data that does not impact individuals' privacy rights or otherwise create a risk of harm. As repeated government and regulators' reports have stressed, it cannot be a choice between the already routine benefits of AI and the protection of personal data: we must find practical ways of ensuring both.

This report will introduce artificial intelligence and some of the technologies enabled by it, as well as some of the challenges and tensions between artificial intelligence and existing data protection laws and principles. The challenges to data protection presented by AI are frequently remarked on but are often addressed only at a surface level. There is an urgent need for a more nuanced, detailed understanding of the opportunities and the issues presented by AI and of practical ways of addressing these challenges, in terms of both legal compliance and ethical issues that AI raises.

We will address specific responses and solutions to the tensions between AI and data protection laws in a separate report. These will include: (1) practices that many organisations are already using, considering new tools and accountability measures; (2) opportunities for interpreting and applying existing data protection laws to AI without stifling its development; and (3) considerations for future data protection laws that account for the demands of AI and other new technologies.

II. Introduction to Artificial Intelligence

Significant advances in the analytical capacity of modern computers are increasingly challenging data protection laws and norms. Those advances are often described as "artificial intelligence", a term that describes the broad goal of empowering "computer systems to perform tasks that normally require human intelligence, such as visual perception, speech recognition, decision-making, and translation between languages".¹ This one term encompasses a wide variety of

technological innovations, each of which may present distinct challenges to existing data protection requirements.

Most AI in use today involves computer systems that perform discrete tasks—for example, playing games, recognising images or verifying identity—by identifying patterns in large amounts of data. The mathematical concept of AI dates back to the 1950s but has found real-world applications in recent years due to advances in processing power and the vast amounts of digital data available for analysis. As a result, AI almost always is associated with “big data”. However, recent applications of AI, such as the use of AI to defeat CAPTCHA and Google’s AlphaGo Zero that taught itself to play Go at the championship level, have occurred with minimal training data, indicating that big data may not always be linked with AI.

All of the above examples are “**narrow**” AI—AI designed to perform one task or set of tasks. Narrow AI is still complicated. As the *New York Times* noted, even narrow AI tools can be “bafflingly opaque” and “evade understanding because they involve an avalanche of statistical probability”.² This is an obvious challenge both for building confidence in new technologies and for compliance with data protection laws.

More challenging are concerns about **artificial general intelligence**. These are “notional future AI system[s] that exhibit apparently intelligent behaviour at least as advanced as a person across the full range of cognitive tasks”.³ When a system can behave in such a way that an observer could not distinguish it from that of a human—it is said to pass the so-called “Turing Test”, set out by Alan Turing in 1950. Such a capability across a wide range of tasks has not yet been achieved.

The ability of a machine to mimic the human brain has led to developments in the field of “machine learning”, which Stanford University professor Andrew Ng has defined as “the science of getting computers to act without being explicitly programmed”.⁴ Machine learning is a subset of AI that has seen many recent developments.

Collectively, these technologies increasingly describe the reality of modern computing, and nations around the globe, from the United States and Canada to EU member states, Japan, Singapore and Australia, have showcased a commitment to be at the forefront of AI with the announcement of ambitious agendas to promote the development of AI technologies. As the European Commission noted in its recent report, *Artificial Intelligence for Europe*: “Artificial intelligence (AI) is already part of our lives—it is not science fiction. From using a virtual personal assistant to organise our working day, to travelling in a self-driving vehicle, to our phones suggesting songs or restaurants that we might like, AI is a reality”. The report goes on to note the important fact that “[b]eyond making our lives easier, AI is helping us to solve some of the world’s biggest challenges: from treating chronic diseases or reducing fatality rates in traffic accidents to fighting climate change or anticipating cybersecurity threats”.⁵ The commitment to AI is further highlighted by the creation of the EU AI Alliance, a multi-stakeholder forum created to promote discussion of all aspects of AI’s advancement, as well as the AI High Level Expert

Group, which is the steering group for the AI Alliance, tasked with drafting ethical guidelines on AI by the end of 2018.⁶

AI and related technologies are rapidly advancing. “Like the steam engine or electricity in the past, AI is transforming our world, our society and our industry”.⁷ Thus, as the term is used below, AI encompasses narrow AI, which is widely used today and has been used for many years, as well as other digital technologies that are ushering in a future of computers so integrated into daily life that we no longer think of them as computers at all.

III. Capabilities of Artificial Intelligence

While machine learning and AI are often used interchangeably, **machine learning** is more accurately understood as one method to achieve AI. Machine learning uses statistical techniques to give computers the ability to “learn”—to progressively improve the machine’s performance by creating new mathematical algorithms—from large volumes of data without being explicitly programmed. Rather than simply following instructions, as traditional computers do, machine learning makes predictions and recommendations based on patterns detected in training data sets.

Machine learning is the basis of other tools, some of which are described below, and it is widely used today to perform numerous tasks, including fraud detection, email filtering, detecting cyber threats such as network intruders or malicious insiders, recommending books or movies, or providing other services based on past or anomalous behaviour. Machine learning is the technology behind Cue, Toyota’s robotic basketball player that has perfect accuracy shooting a basketball and outperforms NBA greats.⁸

Deep learning is a type of machine learning, inspired by the neural networks of the human brain to process successive layers of information and arrive at a conclusion. Deep learning uses multiple layers of artificial neural networks to simulate human decision-making. This technology is at the heart of many AI applications developed today, and enables technologies such as computer vision, text classification, pattern recognition, speech understanding and predictive recommendations. Deep learning has made it possible to have voice recognition technologies throughout our daily lives—in smartphones, digital assistants, AI-powered home security systems and other smart devices. Deep learning in the entertainment industry has enabled Walt Disney to improve significantly image quality in films, as well as improve predictions and understanding of audience reaction to certain scenes.⁹ Often, deep learning uses larger data sets to create larger models and optimally train those models.

Deep learning has enabled a rise in the technology known as **computer vision**, where machines skilled at image recognition, comparison and pattern identification “see” with equal or far greater acuity than human eyes, and then connect what they see based on previously examined training data. Computer vision has created advances in healthcare, national security, assistive care and other various sectors. For example, in health care, algorithms today are able to assess

the risk of heart disease in patients by analysing blood vessels in a retina scan; detect cancerous tumours by examining CT scans; diagnose pneumonia by examining chest x-rays; and identify adult-onset diabetes by looking for patterns of retina damage.¹⁰

Another application of computer vision is helping visually impaired individuals understand images or better perceive their environment by describing them as text, or helping hearing-impaired individuals communicate by translating spoken words to text on a screen.¹¹ Perhaps the most common day-to-day application of computer vision is facial recognition, which is used for accessibility, as well as to unlock smartphones, tag pictures of friends on social media and search images.¹² Computer vision has also proven its use in sports, as auto racing uses it to improve driver safety; golf uses it to improve player experiences and analysis; and the International Gymnastics Federation plans to incorporate it in the Tokyo Olympics of 2020 to assist judges.¹³

Another form of AI technology, **Natural Language Processing (NLP)**, does exactly as the name suggests—interprets and interacts with real-time dialogue. The goal of NLP, which is often combined with speech recognition technologies, is to interact with individuals through dialogue, either reacting to prompts or providing real-time translation among languages. This technology underpins many customer service transactions, as chatbots are often the first line of service. Microsoft’s AI translator is capable of translating Chinese into English with “accuracy comparable to that of a bilingual person”.¹⁴ Facebook is using unsupervised AI for language translation when training data sets are scarce, such as when translating English to Urdu.¹⁵ Such translators have numerous applications spanning across sectors, geographical boundaries and cultural barriers. Major news media have relied on NLP-based technologies to generate thousands of news, sports and financial stories over the past two years, including more than 500 reports in the *Washington Post* about the 2017 elections.¹⁶ Additionally, the GRE exams used for admission to graduate study in many disciplines are graded today by NLP systems.¹⁷

NLP and computer vision are not the only subsets of AI technologies that are driving important advancements in the field, but these two often underpin other applications of AI. For example, **robotics** combines computer vision, NLP and other technologies to train robots to “interact with the world around it in generalizable and predictable ways, ... facilitate manipulation of objects in interactive environments, and ... interact with people”.¹⁸ Robots are beginning to assist in healthcare, at-home care for the sick or elderly and other assistive purposes. In surgeries, robotics technology helps surgeons achieve greater precision and accuracy.

While AI is often perceived as systems acting autonomously, as is the case with home robotics or self-driving vehicles, most practical applications of AI **augment human intelligence**, serving as helpful resources in various professions and automating routine tasks. AI can augment human intelligence by assisting professionals in decision-making, resource management, safety inspection and time management. For example, AI in hospitals is used to suggest diagnoses and treatments to health professionals. In resource allocation, AI is becoming essential for determining truck or airline routes and managing deployment of law enforcement resources. To

assist safety inspectors, Intel has developed a technology to help oil rig inspectors by using AI to identify and detect bolt corrosion levels and the potential need for replacement. Finally, because AI has proved both efficient and effective at issue-spotting in legal contracts, it is used to assist lawyers, shortening the length of time it takes to perform a task, freeing up time to spend on other tasks and ideally lowering legal costs.¹⁹

Scholars have estimated that as many as one in five workers will have an AI acting as a co-worker by 2022.²⁰ In *Technology Vision 2018*, Accenture identified the “Internet of Thinking”, where humans and machines work hand in hand, describing it as “bringing a new level of technological sophistication to the world”.²¹

IV. Public and Private Uses of Artificial Intelligence

The remarkable developments in AI applications have led to considerable use of AI in the public and private sectors. As noted by the AI report of the UK House of Lords, “AI is a tool which is already deeply embedded in our lives”.²² As a computational tool that can enhance many decision-making processes, AI enables subject-matter experts in every sector to deliver improved services and make unprecedented breakthroughs. AI technologies facilitate commercial interactions and personalised services and products, a trend that is highly demanded by consumers and clients. Personalisation occurs in the private sector through travel management, shopper recommendations and targeted advertising, as well as for societal advancements in medical diagnosis and treatment, personalised education and efficient use of resources. The benefits of AI span across a multitude of sectors, including healthcare, marketing, legal services, automotive, human resources, sustainability, agriculture, entertainment, cybersecurity, law enforcement, military and education. Rather than providing an expansive catalogue of the benefits of AI in each of these sectors, this section will provide an overview of changes in some of the major sectors influenced by emerging AI technologies.

AI in Health and Medicine. AI in healthcare is assisting with research and prevention of diseases as well as diagnosis and treatment of patients. Microsoft’s Project Premonition “aims to detect pathogens before they cause outbreaks—by turning mosquitoes into devices that collect data from animals in the environment”.²³ Microsoft is developing drones that autonomously find mosquito hotspots; deploying robots to collect them; and using “cloud-scale genomics and machine learning algorithms to search for pathogens”.²⁴ Intel’s Collaborative Cancer Cloud is designed to help researchers discover new biomarkers associated with cancer diagnoses and progression.²⁵ In addition to assisting medical research, AI is increasingly used in applications for the practice of medicine—whether that is helping doctors find the right location to operate during surgical procedures or scanning images for early disease detection.²⁶

AI in Transportation. One of the most frequently discussed applications of artificial intelligence is sensor-enabled vehicles. Many modern vehicles include AI technologies that provide assistance when backing up or changing lanes. These tools are found on trains, ships and airplanes as well—almost anything that moves. Wholly autonomous vehicles have also

increasingly become a reality, with more than 10 million miles logged on public streets by driverless vehicles designed to react to changing road conditions and traffic patterns. These sensor-enabled vehicles are transforming transportation and promising dramatic changes in private vehicle ownership and use as well as public transportation.

AI in Financial Services. Within financial services, AI is used to assess the credit of clients, back-test trading models, analyse market impact of trades, interact with customers through chatbots and for regulatory reporting.²⁷ In addition, AI is essential for fraud detection and prevention and is being used by financial service organisations and financial technology firms, including banks, credit card companies and other payment service providers, to combat fraud and financial crimes. It is used widely today to identify patterns of normal and unusual behaviours, spot early indicators of fraud, enable faster and more accurate financial decisions and provide financial service professionals with key information meaningfully integrated from a variety of sources. For example, Mastercard acquired Brighterion in 2017 to incorporate its AI technology for fraud prevention.²⁸ Using this and other technologies, Mastercard has developed algorithms and models using AI to determine the likelihood of whether a transaction is legitimate or fraudulent.

AI in Marketing. AI has proven useful in more efficient and effective marketing, helping companies produce targeted ads to consumers most likely to be interested in specific products (and, conversely, not burdening consumers with ads for products for which they have no interest). For example, Nielsen's Artificial Intelligence Marketing Cloud enables clients to "respond instantly to real-time changes in consumer behavior, resulting in more relevant content and advertising, higher levels of customer engagement and improved ROI".²⁹ Popular technology companies such as Amazon, Netflix, Spotify and Facebook as well as traditional retailers such as Starbucks and Walmart use AI to tailor consumer advertisements and customer experiences.

AI in Agriculture. Agriculture is another area where AI is widely utilised in raising livestock and monitoring crops. Just as the agricultural sector was an early industrial user of GPS, it is an early adopter of AI, finding numerous applications for AI technology. For example, a team of researchers developed AI algorithms to assist small cattle farmers in low-income communities. "These algorithms identify patterns for each animal. This customized analysis is then visualized on a Power BI dashboard ... [machine learning], based on an expert knowledge base, provides actionable recommendations, which are sent to farmers via their mobile phones".³⁰ Other recent AI developments in agriculture focus on monitoring, watering and maintaining crops. For example, IBM's Watson can automatically detect and water small sections of vineyards based on data retrieved via sensors, and this technology is currently being adapted to other crop systems as well.³¹ Other agricultural uses of AI include predicting the effectiveness of fertilizers as well as predicting the performance of hybrid seeds based on the genomic information and identifiers of parent lines, which may also have the potential to aid biomedical research.

AI in Education and Training. Artificial intelligence in education has the ability to transform and individualize the student experience. From an early age, teaching robotics are available to help children learn interactively. Online tutoring companies are using AI to analyse, review and tailor individual learning experiences based on techniques where each student seems most responsive.³² AI in an intelligent tutoring system is able to use machine learning to adapt and respond to students' needs in real time. This could be used to provide tutoring sessions to secondary school students, training sessions for military personnel or various on-the-job trainings. AI is also used today to help with grading exams and preventing plagiarism in student papers and published articles. AI can be used to predict needed skills and help to connect graduates' skills with available job opportunities. For example, Pymetrics, "the Netflix-like recommendation algorithm for jobs", seeks to match individual candidates to companies and jobs based on inferences drawn from data collected during neuroscience games.³³

AI in Cybersecurity. AI is helping organisations to monitor, detect and mitigate cybersecurity threats that increasingly face governments, industry and individuals alike. This is already helping with long-standing cybersecurity issues such as spam filters, malicious file detection and malicious website scanning.³⁴ Alphabet recently released Chronicle, "a cybersecurity intelligence platform that throws massive amounts of storage, processing power, and advanced analytics at cybersecurity data to accelerate the search and discovery of needles in a rapidly growing haystack".³⁵ AI-generated dynamic threat models help predict future attacks.³⁶ AI facilitates more efficient threat monitoring, detection and response.

AI for Public Authorities and Public Services. The potential benefits for AI applications to deliver more efficient government services and to assist public safety and security are expansive and have been implemented at international, national and local levels. Internationally, AI has been combined with drone footage to combat wildlife poaching and illegal logging in Uganda and Malaysia, and these technologies are expected to expand to other countries after achieving promising results.³⁷ AI applications assist law enforcement with fraud detection, traffic control, and algorithms to predict recidivism rates and flight risks. Using predictive crime analytics, AI has helped to efficiently deploy law enforcement.³⁸ AI is helping to identify key people in social networks of Los Angeles, California's homeless youth population to help mitigate the spread of HIV.³⁹

While AI has seen demonstrable and numerous uses to assist public authorities, it also has a substantial capacity to aid in public services such as scientific research or conservation of important monuments. For example, researchers at NASA have partnered with technologists at Intel to develop Automated Crater Detection technology to discover craters, and even water, on the moon. Technologists at Intel are also partnering with the China Foundation for Cultural Heritage Conservation to use drones to build models of deteriorated portions of the Great Wall and use AI to scan these sections to determine the exact number of bricks needed to restore and preserve the Wall.⁴⁰

AI for Data Protection. While some scholars have argued that AI poses a threat to data protection, others have posited that AI can offer opportunities to further bolster it. For example, AI can help companies limit or monitor who is looking at an individual’s data and respond in real-time to prevent inappropriate use or theft of data. Companies are developing AI-based privacy tools, such as privacy bots that remember privacy preferences and try to make them consistent across various sites, and privacy policy scanners that attempt to read and simplify privacy policies so that users can better understand them. Polisis, which stands for “privacy policy analysis”, is an AI that uses machine learning to “read a privacy policy it’s never seen before and extract a readable summary, displayed in a graphic flow chart, of what kind of data a service collects, where that data could be sent, and whether a user can opt out of that collection or sharing”.⁴¹

AI can also be useful in alerting users of suspicious websites, advertisements and other malicious activity. Companies are also using AI to prevent malicious or fake content on their online platforms. For example, Facebook is using AI to monitor and respond to fake accounts and inappropriate content on their online platforms.⁴² Finally, AI is enabling companies to develop technologies that are more protective of user privacy. For example, researchers are attempting to develop machine learning techniques that evaluate encrypted data, thereby enhancing user privacy.

V. The Tension with Data Protection

While AI has enormous benefits for society, it also presents a number of challenges, including potential discrimination, antitrust issues or the impact on labor markets. Each of these important issues requires thoughtful attention, but they are beyond the scope of this report because they are the subject of other bodies of law. This report focuses exclusively on the impact that data protection law may have on AI used today and under development for use in the near future. Our next report will address practical steps for industry and regulators to manage those challenges.

A. AI and the Definition of Personal Data

Data protection laws apply when personal data is involved, although the definition of personal data can vary by jurisdiction and by statute. Furthermore, the line between what is “personal” and what is not has been blurred by the correlations and inferences that can be made from aggregated data sets. Today, information that once seemed to be non-personal now has the potential to be personal data, particularly where distinct data elements are joined together. Data users and regulators alike are faced with the difficult task of determining which data should be the subject of regulation.

The EU General Data Protection Regulation (GDPR) defines personal data as:

any information *relating to* an identified or identifiable natural person ('data subject'); an identifiable natural person is one who *can be identified, directly or indirectly*, in particular by reference to an identifier such as a name, an identification number, location data, an online identifier or to one or more factors specific to the physical, physiological, genetic, mental, economic, cultural or social identity of that natural person.⁴³

Other countries also broadly define personal data. For example, under South Korea's Personal Information Protection Act, personal information means "information pertaining to any living person that makes it possible to identify such individual by their name and resident registration number, image, etc.", and specifically includes "information which, if not by itself, makes it possible to identify any specific individual if combined with other information".⁴⁴

AI, and the variety of data sets on which it often depends, only exacerbates the challenge of determining when data protection laws apply by expanding the capability for linking data or recognising patterns of data that may render non-personal data identifiable. This is not a new discovery. As *The Economist* wrote in 2015, "the ability to compare databases threatens to make a mockery of [data] protections".⁴⁵ Simply stated, the more data available, the harder it is to de-identify it effectively.

AI expands on the ability in some settings to make non-personal data identifiable in two ways. First, it broadens the types of and demand for collected data, for example, from the sensors in cell phones, cars and other devices. Second, it provides increasingly advanced computational capabilities to work with collected data. Facial features, gait, fingerprint and other forms of biometric recognition technologies provide an apt example: this expanded data set of discrete, nearly meaningless data points provides greater opportunities for the data to be combined in a way to reliably identify individuals.

Further complexity exists because personal data may be gathered even though identification of a specific individual may not be necessary for AI to take action and make a decision. For example, the sensors in vehicles might be capable of collecting enough data about pedestrians to identify them, but identification would not be necessary to avoid hitting them. The AI only needs to determine that the object is a pedestrian; any data collected is not meant to identify a specific individual. To provide a second example, to train AI to predict the probability of heart attacks occurring in women over 50, personal health data is needed, but the identification of specific individuals is not required for the AI model's analysis.

Finally, while data protection laws attempt to protect sensitive data and similar variables, technologists would argue that algorithms need to include such data in the analysis to ensure accurate and fair results. Moreover, such data may prove useful for human intervention to review and mitigate discrimination or bias. For example, when predicting the likelihood of death in pneumonia patients, researchers at Microsoft discovered that a history of asthma resulted in a lower risk of death, likely because these individuals would seek earlier treatment.

Prior to conducting the analysis, a non-modifiable risk factor such as asthma may not have been inherently obvious or relevant in determining a lower risk of death in pneumonia patients. It is improbable that a history of asthma actually lowers the risk of death, although an AI that excluded this variable would have suggested so without easily identifying the bias. Because protected variables were left in the model, it was easier for researchers to account for them.

Understanding and resolving the scope of data protection law and principles in the rapidly changing context of AI is not an easy task, but it is essential to avoid burdening AI with unnecessary regulatory requirements or with uncertainty about whether or not regulatory requirements apply. As Singapore's Personal Data Protection Commission wrote in its recent discussion paper on AI: "Governance frameworks around AI should be technology-neutral and 'light-touch'", and should provide "regulatory clarity [for] developing AI technologies and translating them into AI solutions".⁴⁶ Clarifying the application of data protection law is also critical to ensuring that scarce resources are not wasted on protecting data that does not impact individuals' privacy rights or otherwise create a risk of harm to them.

B. Data Protection Principles and Requirements

Most data protection laws reflect long-established principles. The OECD Guidelines on the Protection of Privacy and Transborder Flows of Personal Data, adopted in 1980, articulate eight basic principles of data protection: collection limitation, data quality, purpose specification, use limitation, security safeguards, openness, individual participation and accountability.⁴⁷ Most national data protection laws around the world include requirements based on these principles.

AI is in tension with most of these data protection principles.

Collection Limitation, Purpose Specification and Use Limitation. Most data protection laws require that there be a lawful basis for both collecting and processing data. Under the GDPR, for example, the lawful bases for processing personal data are consent, contractual performance, legal obligation, vital interests, public interests or legitimate interest.⁴⁸ It is unclear what level of detail data protection authorities will require for organisations to demonstrate that they have met a lawful basis for processing. All of these depend on an organisation knowing why the data is collected and how it will be used.

Full knowledge and articulation of purposes for processing is also required by the purpose specification and use limitation principles, which respectively provide that personal data should be collected for specified purposes and then used only for those purposes or for purposes that are compatible with the original purposes.

The challenge, as well as the opportunity, is how to comply with these requirements in the context of AI when training data may potentially yield unforeseen and sometimes unpredictable results, by advanced algorithms that are not always directed by or initially understood by their programmers and may increasingly be created only by computers.

Advancements in AI challenge the collection limitation, purpose specification and use limitation principles, but may be further advanced through the reasonable application of these principles.

Moreover, the volume and variety of data typically involved in the development and deployment of AI are enormous. AI technology can use vast amounts of diverse data to improve itself and its interaction with humans. As the Norwegian Data Protection Authority explained: “Most applications of artificial intelligence require huge volumes of data in order to learn and make intelligent decisions”.⁴⁹ In fact, rather than sample data, AI often works by, in the words of the United Kingdom Information Commissioner, “collecting and analysing *all* of the data that is available”.⁵⁰ Providing the necessary volume and variety of data typically requires using data from different sources, where data may have been collected for a different purpose. Denying access to some or all of that data, whether for data protection or other reasons and whether by substantive limits or transactional burdens, necessarily weakens AI and may introduce some unintended bias, because, as articulated by the Norwegian Data Protection Authority, “AI learns from *all* the data it sees”.⁵¹ If AI systems are trained on a limited dataset, representative of only a small segment of the population, these systems will propagate a biased and narrow point of view.

The potential challenge created by a rigid interpretation of these principles is exacerbated by the fact that the collection limitation, purpose specification and use limitation principles undergird most other elements of modern data protection laws. These principles are the foundation of many other legal requirements, such as the need to be transparent and provide privacy notice to individuals, or the need to obtain informed consent for certain data processing. For example, the GDPR provides that consent “should be given by a clear affirmative act establishing a freely given, specific, informed and unambiguous indication of the data subject’s agreement to the processing of personal data relating to him or her”.⁵² How can consent be “specific, informed and unambiguous” if an organisation may not be fully aware of how the collected data will be used, or of all subsequent purposes of processing at the time of collection? Moreover, how can it be established by a “clear affirmative act” given the volume of data and the number of transactions involved on a daily basis?

Finally, the possible transactional burden imposed by many modern data protection regulations (for example, returning to the individual to obtain new consent for an originally unanticipated use) may slow or block beneficial uses of AI. This is true of both the development and the deployment of AI. AI works at a scale and speed far greater than envisioned by the drafters of many data protection laws. Therefore, the increasing challenge is not just how to fit these modern technologies into regulatory frameworks designed for a different world, but how to do so at a speed and scale necessary to serve the public interest.

Data Minimisation. Implicit in the OECD Guidelines, and made explicit in the GDPR and other modern data protection laws, is another widely shared principle: data minimisation. “Personal data must be adequate, relevant and limited to what is necessary in relation to the purposes for which those data are processed”.⁵³ Indeed, as the Norwegian Data Protection Authority noted

in its report on *Artificial Intelligence and Privacy*: “a controller cannot use more personal data than is necessary, and that the information selected must be relevant to the purpose”.⁵⁴ However, with data seen as the “basic building block of the digital economy”,⁵⁵ the concept of data minimisation—that companies should keep data for as little time as possible, use only the amount and type of data necessary for the model, and only for its specified use—can be seen as counterproductive to developing AI technologies. It is difficult to know in advance “what is necessary” in a world of “surprising correlations” and computer-generated discoveries. The challenges of defining a purpose for processing and only keeping data for that purpose are exacerbated because “it is not possible to predict what the algorithm will learn”, and the “purpose may also be changed as the machine learns and develops”.⁵⁶

If interpreted narrowly, data minimisation, as well as limits on data retention as discussed below, can interfere with effective assessment and oversight of AI. If you restrict access to features such as race and gender, it becomes much more difficult (perhaps impossible) to determine if the AI model is biased on those dimensions. If the model is ultimately determined to be biased, it is more difficult to repair the model if the engineers do not have access to the withheld features. One might think that not allowing access to those features would prevent bias from happening in the first place, but that is not true: usually, there are other features in the data that are being used that correlate one way or another with protected variables like race or gender, and the model will learn to be biased using these correlated features. The bias will now be buried in a sea of unknown correlation, and as a result, will be difficult to detect and repair. Researchers developing facial recognition software have discovered that access to more personal data about people from a wider variety of backgrounds, races and ethnicities improves the accuracy of facial recognition, reduces systemic bias and enhances their ability to demonstrate these to regulators.⁵⁷ These are basic demands of AI services. Therefore, it will be necessary to apply this principle in more flexible and nuanced ways when considering new technologies and their applications.

Retention Limitation. To protect personal data and promote data quality, many data protection frameworks and regulations provide for storage limitation requirements. For example, the GDPR requires that personal data shall be “kept ... for no longer than is necessary for the purposes for which the personal data are processed”, though personal data may be stored longer if it “will be processed solely for archiving purposes in the public interest, scientific or historical research purposes or statistical purposes”.⁵⁸ This limitation is related to the data quality principle, as the French CNIL noted, because “data that is quite simply out of date will lead to errors or malfunctions of varying gravity depending on the sector in question, from the mere dispatch of targeted advertising that does not match my actual profile, to an incorrect medical diagnosis”.⁵⁹ And it implicates also rights of individuals, such as the “right to be forgotten”,⁶⁰ which has found both judicial and regulatory support in Europe and elsewhere, and the right to restriction of processing.⁶¹

The underlying tension is that setting short retention periods and deleting or restricting the use of data after its original purpose has been fulfilled or upon request by an individual would strip organisations and society of the potential benefits of using that data for AI training, deployment and oversight. AI and machine learning technologies allow these models to perform optimally. Yet, keeping data for longer periods or indefinitely may violate current data protection laws in the eyes of regulators.

Transparency. The openness and individual participation principles require that data processing be transparent and that individuals are informed about uses of their personal data. The GDPR demands that controllers describe their data processing in greater detail and with concise, intelligible and easily accessible information. The law specifically requires processing to be transparent and further requires organisations to provide individuals the specifics of data processing, including the logic behind any automated decision-making that has legal effect or a similarly significant impact on individuals.⁶² These can be difficult requirements to meet with respect to decisions made by complex AI algorithms, which are often unanticipated.

Data protection principles of transparency and openness are challenged in AI by what many refer to as the “black box” problem. This phenomenon occurs where, as described by the Norwegian Data Protection Authority, the “advanced technology employed is difficult to understand and explain”, and where the neural networks—or successive layers within the technology—make it “practically impossible to explain how information is correlated and weighted in a specific process”.⁶³ This is particularly concerning due to a fear of algorithmic bias. As the AI Forum of New Zealand explained, “AI systems are fed training data by their creators. If this data contains bias, then clearly the system will learn the same bias”.⁶⁴ Inside a “black box”, detecting and understanding the presence of bias is more difficult.

Providing transparency in light of the “black box” phenomenon has been one of the major topics of AI discussed by policymakers, academics and researchers. As Georgetown professor Paul Ohm has stressed, when a program “thrives on surprising correlations and produces inferences and predictions that defy human understanding ... [h]ow can you provide notice about the unpredictable and unexplainable?”⁶⁵ Moreover, the opacity often found in AI models has led many countries to reaffirm a need for transparency in data protection regimes. The French Commission Nationale de l’Informatique et des Libertés (CNIL) recently articulated a “principle of continued attention and vigilance”, noting that this principle could “offset the phenomenon of excessive trust and weakened accountability which can arise in front of ‘black box’ algorithms”.⁶⁶ Technology companies are also working on ways to develop explainable AI, which would also enhance the transparency principle.

Data Quality, Access and Correction. Another consideration with AI and decision-making is data quality and the need for individuals to be able to identify and correct their data. AI technology, like any data-driven technology, can be hindered by inaccurate, incomplete or non-representative data sets, so by making decisions in a “black box”, accuracy and fairness become

a substantial concern. As Singapore’s Personal Data Protection Commission recently explained in a discussion paper on AI, data accountability and accuracy are impacted by “the completeness of the data required, how recently the data was collected and updated, whether the data is structured in a machine-understandable form, and the source of the data”.⁶⁷ In its Technology Vision 2018, Accenture identifies the trend of “Data Veracity”, explaining that “the potential harm from bad data can become an enterprise level existential threat”.⁶⁸

When decisions are made using AI, it can be challenging to contest given the complexity of an algorithm—even if some of the data points used to make the decision were incorrect. There is, however, incentive for both AI developers and privacy advocates to address this challenge. AI developers would like to have the most accurate data possible to promote trustworthiness in outcomes. Individuals would like to ensure that the algorithm will not produce a negative outcome based on incorrect, incomplete or insufficient data.

C. Automated Decision-Making and Profiling

The GDPR is distinctive among most data protection laws in that it specifically addresses profiling and automated decision-making and imposes special restrictions on certain forms of solely automated decision-making under Article 22.

Profiling is defined as “any form of automated processing of personal data consisting of the use of personal data to evaluate certain personal aspects relating to a natural person, in particular to analyse or predict aspects concerning that natural person’s performance at work, economic situation, health, personal preferences, interests, reliability, behaviour, location or movements”.⁶⁹ All of the GDPR requirements apply to profiling, as they would to any other form of processing. Article 21 of the GDPR, however, specifically mentions profiling with regard to the right to object.

Similarly, all of the GDPR requirements apply to automated decision-making, though special rules exist for solely automated decision-making. Article 22 provides that an individual has the “right not to be subject to a decision based solely on automated processing, including profiling, which produces legal effects concerning him or her or similarly significantly affects him or her”.⁷⁰ Article 22 reflects the risk-based approach of the GDPR and subjects these significant legal or similar decisions to a higher compliance bar. This is driven by a concern for algorithmic bias; a worry of incorrect or unsubstantiated solely automated decisions based on inaccurate or incomplete data; and the need for individuals to have redress and the ability to contest a decision if an algorithm is incorrect or unfair. There is a concern that AI technology can reinforce or reflect human biases when making decisions.⁷¹

The Article 29 Working Party has provided Guidelines on Automated Decision-Making⁷² that interpret Article 22 as a direct prohibition on such automated decision-making absent the existence of one of three exceptions provided by Article 22(2). This interpretation further limits the number of legal bases that can be used for automated decision-making and notably

prevents the use of legitimate interest as a basis for processing when making such automated decisions.

Article 22 of the GDPR also gives rights to individuals to contest the decision and seek human intervention and review.⁷³ Although one attribute of AI in many settings is the ability to act without human intervention, Article 22 exists to address a concern that handing over full decision-making authority to a machine where its decisional output can produce legal or similarly significant consequences for an individual is potentially harmful and dangerous. Article 22 acts to limit such consequences by providing individuals with the right not to be subject to such automated decision-making and with recourse to human intervention under Article 22(3).

However, it is crucial to understand what constitutes a “legal or similarly significant effect” to prevent stifling of AI innovation and operation. The WP29 guidelines reflect this understanding by noting that “only serious impactful effects will be covered by Article 22”.⁷⁴

Furthermore, the WP29 Guidelines highlight how difficult it may be to avoid the tension between AI and automated decision-making. For example, the Guidelines provide that “[c]ontrollers seeking to rely upon consent as a basis for profiling will need to show that data subjects understand exactly what they are consenting to, and remember that consent is not always an appropriate basis for the processing”.⁷⁵ Therefore, consent may not be an acceptable basis, and in cases when it could be, organisations will have to overcome the challenges already noted with providing sufficient information about AI.

Finally, in its guidelines, the WP29 notes that “[w]hilst there can be advantages to retaining data in the case of profiling, since there will be more data for the algorithm to learn from, controllers must comply with the data minimisation principle when they collect personal data and ensure that they retain those personal data for no longer than is necessary for and proportionate to the purposes for which the personal data are processed”.⁷⁶ The inherent challenge is determining when the purpose ends in relation to an AI application. Storing data indefinitely within a profile is inherent to many applications, and one can argue that it is ultimately more advantageous to individuals in the sense that the more data that is taken into account by a profiling algorithm or automated decision-making process, the more accurate the result will be.

VI. Observations

This First Report has highlighted the capabilities and benefits of AI as well as some of the tensions and challenges presented by the interaction between AI and data protection law. From this discussion, six general observations emerge.

1. **Not all AI is the same.** As we have seen, the term AI is applied to myriad technologies and applications, designed to be used in diverse settings with widely varying consequences. While a computer playing chess has a finite (although large) number of

moves to learn, a computer aiding in surgery could use an infinite amount of training data to perform optimally. Additionally, if an algorithm playing chess makes a mistake, the harm is trivial compared to the substantial harm that could result from AI producing an unexpected result during surgery. While it is possible to make high-level observations about AI generally, when it comes to applying laws and ethical principles, specificity about technologies, applications, contexts and consequences does matter.

2. **AI is widely used in society today and is of significant economic and societal value.** AI is not a new or futuristic concept; it is prevalent today and something individuals interact with constantly in mobile devices, vehicles, homes and businesses. Governments and companies around the world are rapidly investing in AI because of the reality of its substantial benefits in health, commerce, trade, public safety and other areas. This does not mean that AI does not present important issues that must be addressed, but rather that now is the time to consider practical data protection compliance. Equally, it would be counterproductive to impose unnecessary barriers to its development or to the vast amounts of data on which it depends.
3. **AI requires substantial amounts of data to perform optimally.** Data is the oxygen of AI. AI requires data to train algorithms and increase accuracy and overall functionality. With few exceptions, more data is better than less, and there is almost never enough. This is necessary not only for AI to achieve its full potential, but also, as we have seen and is described further below, to guard against bias or error and to prevent monopolization of critical AI. As Oxford University Professor Viktor Mayer-Schönberger recently noted in *Foreign Affairs*, even large companies are in need of more data to develop and deploy AI, as “the quality of [AI applications] would deteriorate absent sufficient data, leading to inefficient transactions and reduced consumer welfare”.⁷⁷ This is especially true if AI is to serve the needs of small but vital subsets of the population.
4. **AI requires data to identify and guard against bias.** AI, like the humans who develop it, is not free from bias or error. However, it has the potential to avoid many of the irrational biases that infect human decision-making and to make detecting bias and errors easier and more reliable. However, as we have seen, to do this AI requires access to extensive data, especially including sensitive or protected data. Data on race, ethnicity, gender and other sensitive attributes may assist in the detection and remedy of bias or discrimination in AI (and other) models. Denying access to or preventing retention of such data will only make it harder to detect and remedy bias while also denying all segments of society the full potential of AI’s benefits. At the same time, it is important to carefully control the availability and use of such data to ensure that it is not used to facilitate discrimination.
5. **The role of human oversight of AI is likely to and will need to change for AI to deliver the greatest benefit to humankind.** Society must make intelligent, well-informed and thoughtful decisions about the role of AI, but as the speed, accuracy and impact of AI

increases, the role of human oversight will need to change. Although many current applications of AI are designed to augment human intelligence, in the face of autonomous, rapid AI, human intervention may be not only unnecessary but counterproductive. Human decision-making is sometimes unexplainable or irrational. We should aspire for AI to operate more efficiently and accurately than humans, and to make less biased, more rational and reliable decisions. Individuals may not always understand how specific AI works, but they can assure that it is developed according to legal and ethical principles. Humans are essential to evaluating its results and providing redress in the case of incorrect or unfair decisions.

6. **AI challenges some requirements of data protection law.** Companies can and must strive to comply with data protection law as it currently exists. Given the distinctive characteristics of AI, this will require forward-thinking practices by companies and reasonable interpretation of existing laws by regulators if individuals are to be protected effectively and society is to enjoy the benefits of advanced AI tools. In addition, as new data protection laws are adopted, there will also be an opportunity to consider whether there are more effective approaches to protecting privacy in the context of AI and other new technologies. Many technological advances—the proliferation of mobile devices, the growth of IOT, the advent of big data—have already posed challenges to parts of the OECD Guidelines and the laws based on them. AI is likely to exacerbate those challenges, but it also creates opportunities for including more creative approaches in new laws to ensure that the public enjoys the benefits of advanced technologies while also having confidence that individual privacy is assured.

In CIPL's Second Report on Delivering Sustainable AI Accountability in Practice, we will address some of the critical tools that companies and organisations are starting to develop and implement to promote accountability for their use of AI within existing legal and ethical frameworks, as well as reasonable interpretations of existing principles and laws that regulators can employ to achieve efficient, effective privacy protection in the AI context. Finally, it will discuss considerations for the development of future data protection laws that account for the development of AI and other innovative technologies.

REFERENCES

- ¹ English Oxford Living Dictionaries, “Artificial Intelligence”, available at https://en.oxforddictionaries.com/definition/artificial_intelligence.
- ² Kuang, C., Can A.I. Be Taught to Explain Itself?, New York Times Magazine (21 November 2017), available at https://www.nytimes.com/2017/11/21/magazine/can-ai-be-taught-to-explain-itself.html?_r=0.
- ³ Executive Office of the President of the United States, National Science and Technology Council Committee on Technology, Preparing for the Future of Artificial Intelligence (October 2016), available at https://obamawhitehouse.archives.gov/sites/default/files/whitehouse_files/microsites/ostp/NSTC/preparing_for_the_future_of_ai.pdf.
- ⁴ Machine Learning, Coursera, available at <https://www.coursera.org/learn/machine-learning>.
- ⁵ Communication from the Commission, Artificial Intelligence for Europe, COM (2018) 237 final, available at http://ec.europa.eu/newsroom/dae/document.cfm?doc_id=51625.
- ⁶ High Level Expert Group on Artificial Intelligence, (14 June 2018), available at <https://ec.europa.eu/digital-single-market/en/high-level-expert-group-artificial-intelligence>.
- ⁷ Id.
- ⁸ Camparo, A., This basketball-playing robot is so good it could outshoot Stephen Curry, nbcnews.com (20 March 2018), available at <https://www.nbcnews.com/mach/science/basketball-playing-robot-so-good-it-could-outshoot-stephen-curry-ncna858011>.
- ⁹ Rayo, E., Artificial Intelligence at Disney, Viacom, and Other Entertainment Giants, TechEmergence (11 February 2018), available at <https://www.techemergence.com/ai-at-disney-viacom-and-other-entertainment-giants/>.
- ¹⁰ Timmer, J., AI trained to spot heart disease risks using retina scan, arstechnica.com (24 February 2018), available at <https://arstechnica.com/science/2018/02/ai-trained-to-spot-heart-disease-risks-using-retina-scan/>.
- ¹¹ Seeing AI App, Microsoft Accessibility Blog (12 July 2017), available at <https://blogs.msdn.microsoft.com/accessibility/2017/07/12/seeing-ai-app-is-now-available-in-the-ios-app-store/>; Zee, S., Whose Sign Is It Anyway? AI Translates Sign Language Into Text, blogs.nvidia.com (11 May 2017), available at <https://blogs.nvidia.com/blog/2017/05/11/ai-translates-sign-language/>.
- ¹² Quiñonero Candela, J., Managing Your Identity on Facebook With Face Recognition Technology, Facebook Newsroom (19 December 2017), available at <https://newsroom.fb.com/news/2017/12/managing-your-identity-on-facebook-with-face-recognition-technology/>.
- ¹³ Greenberg, N., PGA Tour Is Embracing Artificial Intelligence, And It Could Change How You Watch Golf, The Roanoke Times (8 July 2018), available at https://www.roanoke.com/washingtonpost/sports/pga-tour-is-embracing-artificial-intelligence-and-it-could-change/article_f46d97b1-0b99-5495-a9e9-a015d0b9620b.html.
- ¹⁴ Del Bello, L., AI Translates News Just as Well as a Human Would, futurism.com (16 March 2018), available at <https://futurism.com/ai-translator-microsoft/>.
- ¹⁵ Johnson, K., Facebook is using unsupervised machine learning for translations, Venture Beat (31 August 2018), available at <https://venturebeat.com/2018/08/31/facebook-is-using-unsupervised-machine-learning-for-translations/>.
- ¹⁶ Keohane, J., What News-Writing Bots Mean for the Future of Journalism, Wired (16 January 2017), available at <https://www.wired.com/2017/02/robots-wrote-this-story/>.

- ¹⁷ Hardesty, L., Is MIT Giving Away the Farm?, MIT Technology Review (21 August 2012), available at <https://www.technologyreview.com/s/428698/is-mit-giving-away-the-farm/>.
- ¹⁸ Stone, P., Brooks, R., Brynjolfsson, E., Calo, R., Etzioni, O., Hager, G., Hirschberg, J., Kalyanakrishnan, S., Kamar, E., Kraus, S., Leyton-Brown, K., Parkes, D., Press, W., Saxenian, A., Shah, J., Tambe, M., and Teller, A., "Artificial Intelligence and Life in 2030". One Hundred Year Study on Artificial Intelligence: Report of the 2015-2016 Study Panel, Stanford University, (September 2016), available at https://ai100.stanford.edu/sites/default/files/ai100report10032016fnl_singles.pdf.
- ¹⁹ Chin, M., An AI just beat top lawyers at their own game, Mashable (26 February 2018), available at <https://mashable-com.cdn.ampproject.org/c/s/mashable.com/2018/02/26/ai-beats-humans-at-contracts.amp>.
- ²⁰ Meister, J., AI Plus Human Intelligence Is The Future Of Work, Forbes.com (11 January 2018), available at <https://www.forbes.com/sites/jeannemeister/2018/01/11/ai-plus-human-intelligence-is-the-future-of-work/#789369cf2bba>.
- ²¹ Accenture Technology Vision 2018 Intelligent Enterprise Unleashed (February 2018), at page 14, available at https://www.accenture.com/t20180227T215953Z_w_us-en/acnmedia/Accenture/next-gen-7/tech-vision-2018/pdf/Accenture-TechVision-2018-Tech-Trends-Report.pdf#zoom=50.
- ²² House of Lords Select Committee in Artificial Intelligence, AI in the UK: Ready, Willing and Able?, HL Paper 100 (2018), available at <https://publications.parliament.uk/pa/ld201719/ldselect/ldai/100/100.pdf>.
- ²³ Project Premonition aims to detect pathogens before they cause outbreaks, Microsoft Project Premonition (Established 2 March 2015), available at <https://www.microsoft.com/en-us/research/project/project-premonition/#>.
- ²⁴ Id.
- ²⁵ Artificial Intelligence, The Public Policy Opportunity, Intel (18 October 2017), available at <https://blogs.intel.com/policy/files/2017/10/Intel-Artificial-Intelligence-Public-Policy-White-Paper-2017.pdf>.
- ²⁶ Project InnerEye – Medical Imaging AI to Empower Clinicians, Microsoft Project InnerEye (Established 7 October 2008), available at <https://www.microsoft.com/en-us/research/project/medical-image-analysis/>; Novartis Cataracts Surgery; Google Retina Scan.
- ²⁷ Financial Stability Board, Artificial Intelligence and Machine Learning in Financial Services, Market Developments and Financial Stability Implications, Report (1 November 2017), available at <http://www.fsb.org/wp-content/uploads/P011117.pdf>.
- ²⁸ Bary, E., Visa and Mastercard Earnings: More Than Just Payments at Play, MarketWatch (25 July 2018), available at <https://www.marketwatch.com/story/visa-and-mastercard-earnings-more-than-just-payments-at-play-2018-07-23>.
- ²⁹ Nielsen Launches Artificial Intelligence Technology, Nielsen.com (4 April 2017), available at <http://www.nielsen.com/us/en/press-room/2017/nielsen-launches-artificial-intelligence-technology.html>.
- ³⁰ Spencer, G., Buffaloes and the Cloud: Students turn to tech to save poor farming families, news.microsoft.com (27 September 2017), available at <https://news.microsoft.com/apac/features/saving-farming-families-tech-one-cow-goat-buffalo-time/>.
- ³¹ Vanian, J., How IBM is Bringing Watson to Wine, fortune.com (9 January 2016), available at <http://fortune.com/2016/01/09/ibm-bringing-watson-wine/>.
- ³² Devlin, H., Could online tutors and artificial intelligence be the future of teaching?, The Guardian (26 December 2016), available at <https://www.theguardian.com/technology/2016/dec/26/could-online-tutors-and-artificial-intelligence-be-the-future-of-teaching>.

- ³³ Hiring Based in Neuroscience + Data Science, Pymetrics, available at <https://www.pymetrics.com/science/>.
- ³⁴ Tully, P., Using defensive AI to strip cyberattackers of their advantage, venturebeat.com (6 March 2018), available at <https://venturebeat.com/2018/03/06/using-defensive-ai-to-strip-cyberattackers-of-their-advantage/>.
- ³⁵ Oltsik, J., Artificial intelligence and cybersecurity: The real deal, csoonline.com (25 January 2018), available at <https://www.csoonline.com/article/3250850/security/artificial-intelligence-and-cybersecurity-the-real-deal.html>.
- ³⁶ Supra note 3.
- ³⁷ Kratochwill, L., Artificial Intelligence Fights Wildlife Poaching, popsci.com (22 April 2016), available at <https://www.popsci.com/national-science-foundation-fights-poaching-with-artificial-intelligence>.
- ³⁸ Rieland, R., Artificial Intelligence Is Now Used to Predict Crime. But Is It Biased?, Smithsonian.com (5 March 2018), available at <https://www.smithsonianmag.com/innovation/artificial-intelligence-is-now-used-predict-crime-is-it-biased-180968337/>.
- ³⁹ Clay, J., USC researcher, and AI, give homeless youth a helping hand with HIV education, USC News (14 July 2017), available at <https://news.usc.edu/124831/usc-researcher-and-ai-give-homeless-youth-a-helping-hand-with-hiv-education/>.
- ⁴⁰ Intel Technology Aids in Preserving the Great Wall of China (16 July 2018), available at <https://newsroom.intel.com/news/intel-technology-aids-preserving-great-wall-china/>.
- ⁴¹ Greenberg, A., An AI That Reads Privacy Policies So That You Don't Have To, wired.com (9 February 2018), available at <https://www.wired.com/story/polisis-ai-reads-privacy-policies-so-you-dont-have-to/>.
- ⁴² For example, Facebook uses machine learning to detect and block millions of fake accounts every day by analyzing certain suspicious behaviours without ever assessing the content itself. The company also uses AI to identify and remove terrorist propaganda before it is ever reported by users. In fact, 99 percent of ISIS and Al Qaeda related terror content is removed from Facebook before the online community has flagged it. Community Standards Enforcement Preliminary Report, Facebook (2018), available at <https://transparency.facebook.com/community-standards-enforcement#terrorist-propaganda>.
- ⁴³ GDPR, article 4(1).
- ⁴⁴ Article 2(1) South Korea Personal Information Protection Act, official English translation available at <http://law.go.kr/engLsSc.do?menuId=0&subMenu=5&query=%EA%B0%9C%EC%9D%B8%EC%A0%95%EB%B3%B4%EB%B3%B4%ED%98%B8%EB%B2%95#>.
- ⁴⁵ We'll See You, Anon., The Economist (13 August 2015), available at <https://www.economist.com/science-and-technology/2015/08/13/well-see-you-anon>.
- ⁴⁶ Singapore Personal Data Protection Commission, "Discussion Paper on Artificial Intelligence (AI) and Personal Data—Fostering Responsible Development and Adoption of AI", (5 June 2018), at pages 2-3, available at <https://www.pdpc.gov.sg/-/media/Files/PDPC/PDF-Files/Resource-for-Organisation/AI/Discussion-Paper-on-AI-and-PD---050618.pdf>.
- ⁴⁷ OECD Revised Guidelines on the Protection of Privacy and Transborder Flows of Personal Data (2013), available at http://oecd.org/sti/ieconomy/oecd_privacy_framework.pdf.
- ⁴⁸ GDPR, article 6(1).
- ⁴⁹ Artificial Intelligence and Privacy, Datatilsynet (Norwegian Data Protection Authority) at page 4 (January 2018), available at <https://www.datatilsynet.no/globalassets/global/english/ai-and-privacy.pdf>.
- ⁵⁰ Big Data, Artificial Intelligence, Machine Learning and Data Protection, United Kingdom Information Commissioner's Office at page 11 (Version 2.2 - 2017) (emphasis added), available at <https://ico.org.uk/media/for-organisations/documents/2013559/big-data-ai-ml-and-data-protection.pdf>.

- ⁵¹ Supra note 49 at page 5 (emphasis added).
- ⁵² GDPR, recital 32.
- ⁵³ GDPR, recital 39; article 5(1)(c).
- ⁵⁴ Supra note 49 at page 18.
- ⁵⁵ Supra note 46 at page 2.
- ⁵⁶ Supra note 49 at page 18.
- ⁵⁷ Roach, J., Microsoft improves facial recognition technology to perform well across all skin tones, genders, Microsoft AI Blog (26 June 2018), available at <https://blogs.microsoft.com/ai/gender-skin-tone-facial-recognition-improvement/>.
- ⁵⁸ GDPR, article 5(1)(e).
- ⁵⁹ Commission Nationale de l'Informatique et des Libertés, "How Can Humans Keep the Upper Hand?: The Ethical Matters Raised by Algorithms and Artificial Intelligence", (December 2017), at page 39, available at https://www.cnil.fr/sites/default/files/atoms/files/cnil_rapport_ai_gb_web.pdf.
- ⁶⁰ GDPR, article 17.
- ⁶¹ GDPR, article 18.
- ⁶² GDPR, article 12 (transparency); article 13 and 14 (notice); and article 22 (right not to be subject to automated decision-making).
- ⁶³ Supra note 49 at page 19.
- ⁶⁴ Artificial Intelligence Forum of New Zealand, Artificial Intelligence: Shaping a Future New Zealand (March 2018), at page 64, available at <http://resources.aiforum.org.nz/AI+Shaping+A+Future+New+Zealand+Report+2018.pdf>.
- ⁶⁵ Ohm, P., "Changing the Rules: General Principles for Data Use and Analysis", Privacy, Big Data, and the Public Good: Frameworks for Engagement at page 100 (2014).
- ⁶⁶ Supra note 59 at page 50.
- ⁶⁷ Supra note 46 at page 9.
- ⁶⁸ Supra note 21 at page 40.
- ⁶⁹ GDPR, article 4(4).
- ⁷⁰ GDPR, article 22.
- ⁷¹ Krigsman, M., Artificial Intelligence and Privacy Engineering: Why It Matters NOW, zdnet.com (18 June 2017), available at <http://www.zdnet.com/article/artificial-intelligence-and-privacy-engineering-why-it-matters-now/>.
- ⁷² Article 29 Data Protection Working Party, WP251 Guidelines on Automated Individual Decision-Making and Profiling for the Purposes of Regulation 2016/679 (last Revised and Adopted on 6 February 2018) at page 19, available at http://ec.europa.eu/newsroom/article29/document.cfm?doc_id=49826.
- ⁷³ GDPR, article 22(3).
- ⁷⁴ Supra note 72 at page 21.
- ⁷⁵ Id.
- ⁷⁶ Id., at page 12.
- ⁷⁷ Mayer-Schönberger, V., and Range, T., "A Big Choice for Big Tech: Share Data or Suffer the Consequences", *Foreign Affairs* (Sept./Oct. 2018), p. 52.