

Artificial Intelligence and Data Protection: Observations on a Growing Conflict

Fred H. Cate & Rachel Dockery¹

Abstract

Artificial intelligence (AI) has rapidly developed in recent years. Today, AI tools are used increasingly by both private and public sector organizations around the globe. 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 the tension between AI and data protection laws. As a result, we have both an opportunity and an obligation to examine the effectiveness of current data protection laws in light of 21st-century technological realities.

While compliance with existing data protection laws is important, a better long-term approach is to see the challenges presented by AI as another wake-up call that our current approach to data protection is increasingly outdated and ineffective. Viewed in this light, it is data protection law that must be improved if it is to protect privacy, effectively address the challenges presented by AI, and avoid creating unnecessary, bureaucratic barriers to AI's benefits.

Five reforms appear necessary:

- Shifting from Individual Consent to Data Stewardship
- A More Systemic and Well-Developed Use of Risk Management
- A Greater Focus on Data Uses and Impacts
- A Framework of Harms
- Transparency and Redress

¹ Fred H. Cate is Vice President for Research, Distinguished Professor, and C. Ben Dutton Professor of Law at Indiana University, and a Global Policy Advisor at the Centre for Information Policy Leadership. Rachel Dockery is a Research Fellow in Cybersecurity Law at Indiana University Maurer School of Law and a Fellow of the IU Center for Applied Cybersecurity Research. Portions of this article are excerpted from the Centre for Information Policy Leadership's white paper, *Artificial Intelligence and Data Protection: Delivering Sustainable AI Accountability in Practice*, available at <https://www.informationpolicycentre.com/>. The authors wish to thank the members and staff of the Centre for their generous support of this research. The authors alone are responsible for the views expressed herein.

I. Introduction

Artificial intelligence (AI) has rapidly developed in recent years. Today, AI tools are used increasingly by both private and public sector organizations around the globe. 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 the tension between AI and data protection laws. As a result, we have both an opportunity and an obligation to examine the effectiveness of current data protection laws in light of 21st-century technological realities. We need data protection laws and practices that protect privacy effectively in an era of AI and the big data on which it often depends, but that also do not impose unnecessary roadblocks for the future development of these innovative technologies. 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 article introduces AI and some of the applications enabled by it, as well as some of the challenges and tensions between AI and existing data protection laws and principles. It seeks to provide a more nuanced understanding of those applications and argues that their interaction with data protection laws necessitates, and provides a welcome opportunity for, revising those laws to reflect 21st-century technological realities.

II. Introduction to Artificial Intelligence

A. Defining Artificial Intelligence

The term “artificial intelligence” (AI) 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.”² This one term encompasses a wide variety of technical innovations, each of which may present distinct challenges to existing data protection tools.

Most AI in use today involves computer systems that perform discrete tasks—for example, playing games, recognizing 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-

² English Oxford Living Dictionaries, “Artificial Intelligence,” available at https://en.oxforddictionaries.com/definition/artificial_intelligence.

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 usually is associated with “big data.”³

We have witnessed many examples of “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.”⁴ More challenging are concerns about “artificial general Intelligence.” These are “notional future AI system[s] that exhibit apparently intelligent behavior at least as advanced as a person across the full range of cognitive tasks.”⁵ 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.

Collectively, these technologies increasingly describe the reality of modern computing, and nations around the globe 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.”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.

³ Some recent applications of AI, such as the use of AI to defeat CAPTCHA and Google’s AlphaGoZeros that taught itself to play Go at the championship level, have occurred with minimal training data, suggesting that AI may not always be linked to big data.

⁴ Kuang, C., “Can A.I. Be Taught to Explain Itself?,” *New York Times Magazine* (21 Nov. 2017), available at https://www.nytimes.com/2017/11/21/magazine/can-ai-be-taught-to-explain-itself.html?_r=0.

⁵ Executive Office of the President of the United States, National Science and Technology Council Committee on Technology, *Preparing for the Future of Artificial Intelligence* (Oct. 2016), available at https://obamawhitehouse.archives.gov/sites/default/files/whitehouse_files/microsites/ostp/NSTC/preparing_for_the_future_of_ai.pdf.

⁶ 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.

⁷ Id.

B. Capabilities of Artificial Intelligence

“Machine learning” is a subset of AI that Stanford University professor Andrew Ng has defined as “the science of getting computers to act without being explicitly programmed.”⁸ While the terms 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 cyberthreats such as network intruders or malicious insiders, recommending books or movies, or providing other services based on past or anomalous behavior. Machine learning is the technology behind Cue, Toyota’s robotic basketball player that has perfect accuracy shooting a basketball and outperforms NBA greats.⁹

Deep learning is a type of machine learning, inspired by the neural networks of the human brain to process hidden layers of information and arrive at a conclusion. Deep learning uses multiple layers of artificial neural networks to simulate decision-making of a human. 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. 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 with previously examined data. Computer vision has created advances in health care, 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 analyzing blood vessels in a retina scan; detect cancerous tumors

⁸ Machine Learning, Coursera, available at <https://www.coursera.org/learn/machine-learning>.

⁹ Camparo, A., This basketball-playing robot is so good it could outshoot Stephen Curry, nbcnews.com (20 Mar. 2018), available at <https://www.nbcnews.com/mach/science/basketball-playing-robot-so-good-it-could-outshoot-stephen-curry-ncna858011>.

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 to unlock smart phones, tag pictures of friends on social media, and search images. 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 as the International Gymnastics Federation plans to incorporate it in the Tokyo Olympics of 2020 to assist judges.¹²

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.”¹³ These 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.¹⁴ Additionally, the GRE exams used for admission to graduate study in many disciplines are graded today by NLP systems.¹⁵

NLP and computer vision are not the only subsets of AI technologies that are driving advancements in the field, but these are often the two that underpin other applications of AI.

¹⁰ Timmer, J., AI trained to spot heart disease risks using retina scan, arstechnica.com (24 Feb. 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 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.

¹³ Del Bello, L., AI Translates News Just as Well as a Human Would, futurism.com (16 Mar. 2018), available at <https://futurism.com/ai-translator-microsoft/>.

¹⁴ Keohane, J., “What News-Writing Bots Mean for the Future of Journalism,” *Wired* (16 Jan. 2017), available at <https://www.wired.com/2017/02/robots-wrote-this-story/>.

¹⁵ Hardesty, L., “Is MIT Giving Away the Farm?,” *MIT Technology Review* (21 Aug. 2012), available at <https://www.technologyreview.com/s/428698/is-mit-giving-away-the-farm/>.

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.”¹⁶ Robots are beginning to assist in health care, 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 protect against corrosion 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.¹⁷ AI is also used to help judges calculate criminal sentences. Scholars have estimated that as many as one in five workers will have an AI acting as a coworker by 2022.¹⁸

C. Public and Private Uses of Artificial Intelligence

The remarkable developments in AI applications have led to considerable use of AI in public and private sectors. As the UK House of Lords noted in its recent AI report, “AI is a tool which is already deeply embedded in our lives.”¹⁹ As a computational tool that can enhance any decision-making process, AI enables subject matter experts in every sector to deliver improved services and make unprecedented breakthroughs. AI technologies facilitate commercial interactions and personalized services and products, a trend that is highly demanded by consumers and citizens. Personalization occurs in the private sector through travel

¹⁶ Stone, P., et al, "Artificial Intelligence and Life in 2030." *One Hundred Year Study on Artificial Intelligence: Report of the 2015-2016 Study Panel*, Stanford University (Sep. 2016) https://ai100.stanford.edu/sites/default/files/ai100report10032016fnl_singles.pdf.

¹⁷ Chin, M., An AI just beat top lawyers at their own game, Mashable (26 Feb. 2018), available at <https://mashable-com.cdn.ampproject.org/c/s/mashable.com/2018/02/26/ai-beats-humans-at-contracts.amp>.

¹⁸ Meister, J., AI Plus Human Intelligence Is The Future of Work, *Forbes* (11 Jan. 2018), available at <https://www.forbes.com/sites/jeannemeister/2018/01/11/ai-plus-human-intelligence-is-the-future-of-work/#789369cf2bba>.

¹⁹ 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>.

management, shopper recommendations, and targeted advertising, as well as for societal advancements in medical diagnosis and treatment, personalized education, and efficient use of resources. The benefits of AI span across a multitude of sectors, some of which are described below.

- **AI in Health and Medicine**—AI in health care is assisting with research and prevention of diseases as well as diagnosis and treatment of patients. Intel’s Collaborative Cancer Cloud is designed to help researchers discover new biomarkers associated with cancer diagnoses and progression.²⁰ AI is increasingly used for applications in 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.²¹ AI-equipped “robots can analyze data from pre-op medical records to guide a surgeon's instrument during surgery, which can lead to a 21% reduction in a patient's hospital stay.”²² A partnership between the Cleveland Clinic and IBM uses IBM’s Watson to mine big data and help physicians develop more effective and personalized treatment plans.²³ 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.”²⁴ 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.”²⁵
- **AI in Transportation**—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 vehicle safety, private vehicle ownership, and public transportation.
- **AI in Financial Services**—AI is essential for fraud detection and prevention and is being used by financial service organizations and financial technology firms, including banks,

²⁰Artificial Intelligence, The Public Policy Opportunity, Intel (18 Oct. 2017), available at <https://blogs.intel.com/policy/files/2017/10/Intel-Artificial-Intelligence-Public-Policy-White-Paper-2017.pdf>.

²¹ Project InnerEye—Medical Imaging AI to Empower Clinicians, Microsoft Project InnerEye (7 Oct. 2008), available at <https://www.microsoft.com/en-us/research/project/medical-image-analysis/>.

²² Marr, B., “How Is AI Used In Healthcare—5 Powerful Real-World Examples That Show The Latest Advances,” Forbes (27 Jun. 2018), available at <https://www.forbes.com/sites/bernardmarr/2018/07/27/how-is-ai-used-in-healthcare-5-powerful-real-world-examples-that-show-the-latest-advances/#5b20b9975dfb>.

²³ Id.

²⁴ Project Premonition aims to detect pathogens before they cause outbreaks, Microsoft Project Premonition (2 Mar. 2015), available at <https://www.microsoft.com/en-us/research/project/project-premonition/#>.

²⁵ Id.

credit card, and other payment service providers, to combat fraud and financial crime. It is used widely today to identify patterns of normal and unusual behaviors, 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.

- **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). Popular technology companies such as Amazon, Netflix, and Spotify, as well as traditional retailers such as Starbucks use AI to tailor consumer advertisements and customer experiences.
- **AI in Agriculture**—The agricultural sector was an early industrial user of AI, finding numerous applications for AI applications. For example, a team of researchers partnered with Microsoft to develop algorithms that assist cattle farmers by identifying and analyzing patterns for each animal.²⁶ 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.²⁷ 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.
- **AI in Education and Training**—AI is increasingly in education and training. From an early age, teaching robotics are available to help children learn interactively. Online tutoring companies are using AI to analyze, review, and tailor individual learning experiences based on techniques where each student seems most responsive.²⁸ AI in an intelligent tutoring system is able to use machine learning to adapt and respond to students’ needs in real time. AI is also used today to help with grading exams and preventing plagiarism. AI can be used to predict needed skills and help to connect those with appropriate skills with available jobs opportunities. For example, Pymetrics, “the Netflix-like recommendation algorithm for jobs,” seeks to match individual candidates to

²⁶ Spencer, G., *Buffaloes and the Cloud: Students turn to tech to save poor farming families*, news.microsoft.com (27 Sep. 2017), available at <https://news.microsoft.com/apac/features/saving-farming-families-tech-one-cow-goat-buffalo-time/>.

²⁷ Vanian, J., “How IBM is Bringing Watson to Wine,” *Fortune* (9 Jan. 2016), available at <http://fortune.com/2016/01/09/ibm-bringing-watson-wine/>.

²⁸ Devlin, H., “Could online tutors and artificial intelligence be the future of teaching?,” *The Guardian* (26 Dec. 2016), available at <https://www.theguardian.com/technology/2016/dec/26/could-online-tutors-and-artificial-intelligence-be-the-future-of-teaching>.

companies and jobs based on inferences drawn from data collected during neuroscience games.²⁹

- AI in Cybersecurity—AI is helping organizations to monitor, detect, and mitigate the 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.³⁰ 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.”³¹ AI-generated dynamic threat models help predict future attacks.³²
- AI for Public Authorities and Public Services—AI applications are routinely used to deliver more efficient government services and to assist public safety and security. AI has been combined with drone footage to combat wildlife poaching and illegal logging.³³ AI applications assist law enforcement with fraud detection, traffic control, and algorithms to predict recidivism and flight risks. Using predictive crime analytics, AI has been helped to efficiently deploy law enforcement to areas where crimes are more likely to occur at certain times.³⁴ 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.³⁵ AI is also assisting with public services such as public health, scientific research, and resource conservation. 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

²⁹ Hiring Based in Neuroscience + Data Science, Pymetrics, available at <https://www.pymetrics.com/science/>.

³⁰ Tully, P., Using defensive AI to strip cyberattackers of their advantage, venturebeat.com (6 Mar. 2018), available at <https://venturebeat.com/2018/03/06/using-defensive-ai-to-strip-cyberattackers-of-their-advantage/>.

³¹ Oltsik, J., Artificial intelligence and cybersecurity: The real deal, csoonline.com (25 Jan. 2018), available at <https://www.csoonline.com/article/3250850/security/artificial-intelligence-and-cybersecurity-the-real-deal.html>.

³² *Preparing for the Future of Artificial Intelligence*, supra.

³³ 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>.

³⁴ Rieland, R., “Artificial Intelligence Is Now Used to Predict Crime. But Is It Biased?,” *Smithsonian* (5 Mar. 2018), available at <https://www.smithsonianmag.com/innovation/artificial-intelligence-is-now-used-predict-crime-is-it-biased-180968337/>.

³⁵ 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/>.

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.³⁶

- 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, which remember privacy preferences and try to make them consistent across various sites, and privacy policy scanners, which attempt to read and simplify privacy policies for users to more easily understand. 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.”³⁷ AI is also being used to alert users of suspicious websites, advertisements, and other malicious activity. 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..

III. The Challenge for Data Protection

AI presents challenges as well as benefits. While it is already assisting workers in many professions, AI likely will reduce the need for workers in others. It may introduce bias and new forms of discrimination, especially if the data used in AI development only represents partial segments of the population or reflects existing societal bias. AI will likely challenge traditional notions of urban and residential planning, which have large spaces dedicated to parking lots and garages. AI may also raise important antitrust issues, particularly if the data necessary for its development is concentrated in the hands of a few entities. Each of these important issues require thoughtful attention, but they are beyond the scope of this article and in most cases they are the subject of other bodies of law. This article focuses exclusively on data protection challenges presented by AI used today and under development for use in the near future.

³⁶ 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/>.

³⁷ Greenberg, A., “An AI That Reads Privacy Policies So That You Don't Have To,” *Wired* (9 Feb. 2018), available at <https://www.wired.com/story/polisis-ai-reads-privacy-policies-so-you-dont-have-to/>.

A. The Scope of Data Protection Regulation in the AI Context: Personal Data

Data protection laws apply when personal data is involved. Unfortunately, the line between what is “personal” and what is not has been substantially blurred by the correlations and inferences that can be made from aggregated data sets. Information that once seemed to be non-personal now has the potential to be personal data, and 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.³⁸

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.”³⁹

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 recognizing patterns of data that may render non-personal data identifiable. This is not a new discovery. Professor Latanya Sweeney demonstrated that 87% of the US population is uniquely identified with just three data elements: date of birth, gender, and 5-digit ZIP Code.⁴⁰ There are well-publicized examples of Google, Netflix, AOL, and others releasing deidentified data sets only to have the data reidentified within days by researchers correlating them with other data

³⁸ GDPR, supra article 4(1) (emphasis added).

³⁹ 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#>.

⁴⁰ Sweeney, L., Simple Demographics Often Identify People Uniquely, Carnegie Mellon University, Data Privacy Working Paper at 3 (2000), available at <https://dataprivacylab.org/projects/identifiability/paper1.pdf>.

sets.⁴¹ As *The Economist* wrote in 2015, “the ability to compare databases threatens to make a mockery of [data] protections.”⁴²

Even wholly non-identifiable data may act to identify unique users or machines. For example, browser choice and font size can provide an accurate, unique online identifier.⁴³ Simply stated, the more data are available, the harder it is to de-identify them effectively. AI only makes de-identification harder, in two ways. First, it facilitates the demand for more 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: they collect thousands of discrete, nearly meaningless data points and then combine them in a way to provide reliable identification of individuals.

Further complicating the role of personal data, identification 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 personal data collected is not meant to identify a specific individual. Similarly, for AI to predict the probability of heart attacks occurring in women over 50, personal data is needed, but identification of individuals is not.

While data protection laws and regulations attempt to protect sensitive data and similar variables, AI algorithms need to include such data in the analysis to ensure accurate and fair results. 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 are likely to seek earlier treatment. Because those protected variables were left in the model, it was easier for researchers to account for them.

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 protect privacy effectively in this increasingly critical context and to avoid burdening AI with unnecessary regulatory requirements or with uncertainty about whether or not regulatory requirements apply. Clarifying the application of data protection law is also critical to ensuring that scarce resources are not wasted on

⁴¹ Dwork, C., “Differential Privacy: A Cryptographic Approach to Private Data Analysis,” *Privacy, Big Data, and the Public Good: Frameworks for Engagement* at 296 (2014).

⁴² “We’ll See You, Anon,” *The Economist* (13 Aug. 2015), available at <https://www.economist.com/science-and-technology/2015/08/13/well-see-you-anon>.

⁴³ Kirk, J., “EFF: Browsers Can Leave a Unique Trail on the Web,” *PC World* (29 Jan. 2010), available at <https://www.pcworld.com/article/188134/article.html>.

protecting non-personal data that does not impact data protection, nor does it create risks and harms to individuals.

B. Data Protection Principles and Requirements Applied to AI

Most data protection laws reflect principles established in 1980 in the OECD Guidelines on the Protection of Privacy and Transborder Flows of Personal Data. The Guidelines articulate eight basic principles of data protection: collection limitation, data quality, purpose specification, use limitation, security safeguards, openness, individual participation, and accountability.⁴⁴

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

1. Collection Limitation, Purpose Specification, and Use Limitation.

Most data protection laws require that there be a lawful basis for collecting 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.⁴⁵ The question remains how organizations can give data protection authorities confidence that they have considered a lawful basis for processing while still allowing flexibility to AI models. All of these depend on an organization 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, of course, is how to comply with these requirements in the context of AI when data is being used for unforeseen and often unpredictable purposes, by advanced algorithms that are not always understood by their programmers and will increasingly be created only by computers. 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?”⁴⁶

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:

⁴⁴ 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.

⁴⁵ GDPR, supra art. 6(1).

⁴⁶ Ohm, P., “Changing the Rules: General Principles for Data Use and Analysis,” *Privacy, Big Data, and the Public Good: Frameworks for Engagement* at 100 (2014).

“Most applications of artificial intelligence require huge volumes of data in order to learn and make intelligent decisions.”⁴⁷ 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.”⁴⁸ 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 unintended bias.

The challenge is exacerbated by the fact that the collection limitation, purpose specification, and use limitation principles undergird most other elements of modern data protection law, 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.”⁴⁹ How can consent be “specific, informed and unambiguous” if an organization is not fully aware of how the collected data will be used, or of all subsequent purposes of processing? Moreover, how can it be established by a “clear affirmative act” given the volume of data and the number of transactions involved?

Another basis for processing personal data is “legitimate interest,” defined in Article 6 of the GDPR as: “processing [that] is necessary for the purposes of the legitimate interests pursued by the controller or by a third party, except where such interests are overridden by the interests of fundamental rights and freedoms of the data subject which require protection of personal data.”⁵⁰ However, assessing the “purposes of the legitimate interests” and balancing those with the “fundamental rights and freedoms of the data subject” require knowledge of how the data will be used and for what purposes at each stage of the processing. This may be difficult to ascertain in the context of rapidly evolving AI.

The 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

⁴⁷ Artificial Intelligence and Privacy, Datatilsynet (Norwegian Data Protection Authority) at 4 (Jan. 2018), available at <https://www.datatilsynet.no/globalassets/global/english/ai-and-privacy.pdf>.

⁴⁸ Big Data, Artificial Intelligence, Machine Learning and Data Protection, United Kingdom Information Commissioner’s Office at 11 (Version 2.2 - 2017) (emphasis added), at <https://ico.org.uk/media/for-organisations/documents/2013559/big-data-ai-ml-and-data-protection.pdf>.

⁴⁹ GDPR, supra recital 32.

⁵⁰ GDPR, supra art. 6(1)(f)

regulatory frameworks designed for a different world, but how to do so at a speed and scale necessary to serve the public interest.

2. Data Minimization

Implicit in the OECD Guidelines, and made explicit in the GDPR and other modern data protection laws, is another widely shared principle: data minimization. “Personal data must be adequate, relevant and limited to what is necessary in relation to the purposes for which those data are processed.”⁵¹ Indeed, as the Norwegian Data Protection Authority noted in its recent 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.”⁵² However, with data seen as the “basic building block of the digital economy,”⁵³ the concept of data minimization—that companies should keep data for as little time as possible and only for its specified use—is 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.”⁵⁴

3. Transparency

The openness and individual participation principles require that data processing be transparent and that individuals are informed about the uses of their personal data. Because decisions made by AI applications often have complex algorithms that cannot be fully understood or explained, these applications can be in tension with the transparency principle of data protection. 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 organizations 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.⁵⁵

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 hidden layers within the

⁵¹ GDPR, supra recital 39; art. 5(1)(c).

⁵² *Artificial Intelligence and Privacy*, supra at 18.

⁵³ *Id.* at 2.

⁵⁴ *Id.* at 18.

⁵⁵ GDPR, supra art. 12 (transparency); arts. 13 and 14 (notice); and art. 22 (right not to be subject to automated decision-making).

technology—make it “practically impossible to explain how information is correlated and weighted in a specific process.”⁵⁶

4. Data Quality, Access, and Correction

As the challenge for AI and transparency suggests, another concern with AI 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 non-transparent “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.”⁵⁷ The volume and variety of data used in developing and operating most AI applications make compliance with the data quality, access, and correction principles more difficult.

5. Retention Limitation

To protect personal data and promote data quality, many data protection laws provide for storage limitation requirements. For example, the GDPR storage requirements permit “identification of data subjects 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 purpose specification principle, as the right to store data ends when the purpose, and thus the user consent, no longer exists. It also relates to the data quality principle.

The underlying tension is that setting short retention periods and deleting data after its original purpose has been fulfilled would deny individuals, organizations, and society of the potential benefits of using that data for AI training and deployment. Moreover, retaining outdated and even sensitive data can help reduce and reveal bias in AI applications.

⁵⁶ Artificial Intelligence and Privacy, *supra*.

⁵⁷ 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 9, available at <https://www.pdpc.gov.sg/-/media/Files/PDPC/PDF-Files/Resource-for-Organisation/AI/Discussion-Paper-on-AI-and-PD--050618.pdf>.

⁵⁸ GDPR, *supra* art. 5(1)(e).

6. 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.”⁵⁹ 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.”⁶⁰ Article 22 reflects the risk-based approach of the GDPR and subjects these significant legal or similar decisions to a higher compliance bar.⁶¹

The Article 29 Working Party has provided Guidelines on Automated Decision-Making⁶² 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. This threatens to impede use of AI.

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.”⁶³ Therefore, consent may not be an acceptable

⁵⁹ Id., art. 4(4).

⁶⁰ Id., art. 22.

⁶¹ 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/>.

⁶² 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 Feb. 2018) at 19, available at http://ec.europa.eu/newsroom/article29/document.cfm?doc_id=49826.

⁶³ Id.

basis, and in cases when it could be, organizations 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.”⁶⁴ 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.

IV. Rethinking Data Protection

Many regulators, businesses, attorneys, and academics are working hard to find ways to address the challenges presented by AI to data protection laws. These are important initiatives and obviously necessary in light of the urgent need for users of data to comply with existing data protection laws. However, as we have seen, the tension between those laws and AI is so great and so fundamental that efforts to reconcile them run the risk of weakening data protection or interfering with the benefits of AI. Neither result is desirable given the importance of AI and of personal privacy.

A better long-term approach is to see the challenges presented by AI as another wake-up call that our current approach to data protection is increasingly outdated and ineffective. Viewed in this light, it is data protection law that must be improved if it is to protect privacy, effectively address the challenges presented by AI, and avoid creating unnecessary, bureaucratic barriers to AI’s benefits.

Over the past decade there has been considerable attention given to how data protection law might be modernized to work better not only in the face of AI, but the growth of big data, the Internet of Things, social media, and other phenomenon that were not anticipated when the OECD Guidelines were published in 1980.⁶⁵ Much of this work is relevant to the challenges presented by AI. In particular, five reforms appear necessary.

⁶⁴ Id. at 12.

⁶⁵ See, e.g., Cate, F., “Big Data, Consent, and the Future of Data Protection,” in Sugimoto, C., Hamid R. Ekbia & Michael Mattioli, eds., *Big Data is Not a Monolith 2* (MIT 2016); Cate, F., “Protecting Privacy in Big Data,” *Journal of Law and Economic Regulation*, vol. 8, no. 1 (2015); Cate, F., P. Cullen & V. Mayer-Schönberger, *Data Protection for the 21st Century World* (Microsoft 2014); Cate, F. & V. Mayer-Schönberger, *Data Use and Impact*

A. Shifting from Individual Consent to Data Stewardship

Most data protection laws place some or all of the responsibility for protecting privacy on individual data subjects through the operation of notice and consent. Effective governance of the large data sets used by most AI applications requires shifting more responsibility away from individuals and toward data collectors and data users, who should be held accountable for how they use and manage the data.

While specific methods for accomplishing this are discussed below, the recognition that processors should be liable for reasonably foreseeable harms will create a significant incentive for greater care in their collection and use of data, whether in AI or in other ways. It will also restrict the practice of allowing processors to escape responsibility by providing notice and obtaining (or inferring) consent. This should thus reduce the burden imposed on individuals and focus their attention on data processing activities only where there are meaningful, effective choices to be made. Processors, in turn, will benefit by not wasting resources on, or having the development or deployment of AI restricted by, efforts to comply with ineffective measures, such as notices that no one reads or adhering to terms of consent that are often illusory at best.

B. A More Systemic and Well-Developed Use of Risk Management

Risk management is the process of systematically identifying harms and benefits that could result from an activity. Risk management does not alter rights or obligations, but by assessing both the likelihood and severity of harms and benefits, risk management helps organizations identify mitigation strategies and facilitates an optimum outcome that maximizes potential benefits while reducing the risk of harms to that it falls within acceptable limits.⁶⁶

Data protection has long relied on risk management as a tool for complying with legal requirements and ensuring that data are processed appropriately and that the fundamental rights and interests of individuals are protected effectively. Yet these risk management processes, whether undertaken by businesses or regulators, have often been informal and unstructured and failed to take advantage of many of the widely accepted principles and tools of risk management in other areas.

Global Workshop (Center for Applied Cybersecurity Research 2013); Cate, F., P. Cullen & V. Mayer-Schönberger, *Data Protection Principles for the 21st Century* (Oxford Internet Institute 2013), available at <https://www.repository.law.indiana.edu/facbooks/23/>.

⁶⁶ International Organization for Standardization, *ISO 31000:2009 Risk management—Principles and guidelines*. See generally Centre for Information Policy Leadership, *A Risk-based Approach to Privacy: Improving Effectiveness in Practice* (2014), available at http://www.hunton.com/files/upload/Post-Paris_Risk_Paper_June_2014.pdf; Centre for Information Policy Leadership, *The Role of Risk Management in Data Protection* (2014), available at http://www.hunton.com/files/Uploads/Documents/Centre/Role_of_Risk_Management_in_Data_Protection.pdf.

In addition, institutional risk management in the field of data protection has suffered from the absence of any consensus on the harms that risk management is intended to identify and mitigate in the area of data protection. This is the starting point for effective risk assessment in other fields. As a result, despite many examples of specific applications, a risk-based approach still does not yet provide a broad foundation for data protection practice or law.

It is critical that risk management around data protection, while remaining flexible, not continue in the largely ad hoc, colloquial terms in which it has evolved today. In other areas—for example, financial and environmental risk—we have seen the development of a professional practice of risk management, including specialized research, international and sectoral standards, a common vocabulary, and agreed upon principles and processes. The same is needed in data protection risk management. This is especially true with AI, which may present a range of possible risks—some related to data protection, some not. A consistent widely-shared approach to which risks are within the proper purview of data protection law is urgently needed.

There is some movement in this direction already. In 2013 the Council of Ministers of the Organization for Economic Co-operation and Development revised the OECD Guidelines to “implement a risk-based approach.”⁶⁷ In the accompanying Explanatory Memorandum, the drafters noted the “importance of risk assessment in the development of policies and safeguards to protect privacy.”⁶⁸ Some portions of the recently enacted GDPR also reflect this approach.

Risk management holds special promise in the world of AI by facilitating thoughtful, informed decision-making by data collectors and users that takes into account not only their risks but those of data subjects, by explicitly considering both harms and benefits, and by focusing increasingly scarce resources of both data processors and government regulators where they are needed most.

As the editors of Oxford University Press’ *International Data Privacy Law* have opined:

[We] applaud the attention being given to risk management and its role in data protection. In its proper place, risk management can help prioritize the investment of scarce resources in protecting privacy and enforcing privacy

⁶⁷ Organisation for Economic Co-operation and Development, *Supplementary Explanatory Memorandum to the Revised Recommendation of the Council Concerning Guidelines Governing the Protection of Privacy and Transborder Flows of Personal Data* (2013), 30, available at <http://www.oecd.org/sti/ieconomy/2013-oecd-privacy-guidelines.pdf>.

⁶⁸ Organisation for Economic Co-operation and Development, *OECD Guidelines Governing the Protection of Privacy and Transborder Flows of Personal Data*, C(80)58/FINAL, as amended by C92013)79 (2013), 12, available at <http://www.oecd.org/sti/ieconomy/2013-oecd-privacy-guidelines.pdf>.

obligations. It can identify serious risks to privacy and measures for mitigating them. It can expand our collective thinking about the range of risks that the processing of personal data can present to individuals, organizations, and society, especially in a world of nearly ubiquitous surveillance, big data, cloud computing, and an onslaught of Internet-connected devices. And it can help bring rigor and discipline to our thinking about data processing and how to maximize its benefits while reducing its costs.⁶⁹

C. A Greater Focus on Data Uses and Impacts

There is often a compelling reason for personal data to be disclosed, collected, or created. Assessing the risk to individuals posed by those data almost always requires knowing the context in which they will be used. Data used in one context or for one purpose or subject to one set of protections may be both beneficial and desirable, where the same data used in a different context or for another purpose or without appropriate protections may be both dangerous and undesirable.⁷⁰ As a result, data protection should, in the words of the U.S. President’s Council of Advisors on Science and Technology, “focus more on the actual uses of big data and less on its collection and analysis.”⁷¹

Focusing on the use and impact of personal data does not eliminate responsibilities or regulation relating to data collection, nor should a focus on consent in specific or sensitive circumstances be abandoned. Rather, in many situations, a more practical, as well as sensitive, balancing of valuable data flows and more effective privacy protection is likely to be obtained by focusing more attention on appropriate, accountable use.

Under a more use-based approach, data users would evaluate the appropriateness of an intended use of personal data not by focusing primarily on the terms under which the data were originally collected, but rather on the likely impacts of a proposed use and the risks they pose for individuals. Such a focus on use is more intuitive because most individuals and institutions already think about uses when evaluating their comfort with proposed data processing activities. “What are you going to do with the data?” “How do you intend to use it?” “What are the benefits and risks of the proposed use?” These are the types of questions that many individuals ask—explicitly or implicitly—when they inquire about data processing activities. They can be answered only in connection with specific uses or categories of uses, and they are precisely the questions that data users would be required to ask—and answer—regarding proposed uses of data.

⁶⁹ Kuner, C., et al, “Risk Management in Data Protection,” *International Data Privacy Law* 5, no. 2 (2015), 95, available at <https://academic.oup.com/idpl/article/5/2/95/645238>.

⁷⁰ See Nissenbaum, H., *Privacy in Context* (Stanford University Press 2010).

⁷¹ Executive Office of the President, *Big Data: Seizing Opportunities, Preserving Values* xiii (2014), available at https://obamawhitehouse.archives.gov/sites/default/files/docs/big_data_privacy_report_may_1_2014.pdf.

One of the most pronounced changes that will result from the evolution toward a greater focus on use is to diminish the role of the purpose for which data were originally collected. The purpose specification principle is already problematic today for many reasons, including the fact that, precisely because of it, data processors usually specify exceptionally broad purposes that provide little meaningful limit on their subsequent use of data. In addition, because data increasingly are generated in ways that involve no direct contact with the individual (for example, collected by sensors or inferred from existing data) so there is never a purpose specified. Moreover, with the growing use of AI, personal data may have substantial valuable uses that were wholly unanticipated when the data were collected, yet the data were collected in such a way or are so vast as to make contacting each individual to obtain consent for the new use impractical, as well as potentially undesirable where the beneficial use depends on having a complete data set.

Some modern data protection systems have dealt with these problems by creating broad exceptions to this principle, interpreting “not incompatible” so broadly as to undermine the principle, or simply ignoring it altogether. Taking maximum advantage of AI will require a more thoughtful approach to purpose specification. The principle will have less relevance in many settings. This will certainly be true when data are observed or inferred without any contact with the individual, but it will also likely be true in many other settings. Instead, it is the analysis of risks associated with an intended use that determines whether, and subject to what protections, a particular AI application is appropriate.

The terms under which data were collected would remain relevant when a specific purpose is provided at time of collection and is a meaningful factor in obtaining access to data. This would be especially clear in settings where users had made meaningful choices (e.g., specifying a preferred medium for future communications) or where the data processor had agreed to specific limits as a condition of obtaining personal information (e.g., an explicit promise not to share the data).

But a greater focus on risk-assessment of specific uses, rather than focusing on consent, is essential to ensure that data users do not evade their commitments, that valuable uses of data are not inappropriately deterred, and that data protection laws are not claimed to reflect a principle that increasingly they do not. As Professor Susan Landau wrote in 2015 in *Science*: data protection laws have attempted to protect privacy “through notice and consent. But for reasons of complexity (too many tiny collections, too many repurposings) those are no longer effective. . . . [T]he value of big data means we must directly control use rather than using notice and consent as proxies.”⁷²

A use-focused approach is especially important in the context of AI because of the difficulty—in some cases, impossibility—of explaining in advance and in lay terms which data will be needed and for precisely which purposes. As a result, data protection based on a notice specifying

⁷² Landau, S., “Control use of data to protect privacy,” *Science* 347:6221, Jan. 30, 2015, at 504.

intended uses of data and consent for collection based on that notice can result in blocking socially valuable uses of data, lead to meaninglessly broad notices, or require exceptions to the terms under which the individual consented. If privacy protection is instead based on a risk analysis of a proposed use and its outcome, then it is possible to achieve an optimum benefit from the use of the data and optimum protection for data fine-tuned for each intended use.

D. A Framework of Harms

Measuring risks connected with data uses is especially challenging because of the intangible and subjective nature of many perceived harms. Any risk assessment must be both sufficiently broad to take into account the wide range of harms (and benefits), and sufficiently simple, so that it can be applied routinely and consistently. Perhaps most importantly, the assessment should be transparent to facilitate fairness, trust, and future refinement.

The goal of a risk management approach focused on data uses is to reduce or eliminate the harm that personal information can cause to individuals. Accomplishing this, however, requires a clear understanding of what constitutes “harm” or other undesired impact in the privacy context. Surprisingly, despite almost 50 years of experience with data protection regulation, that clear understanding is still lacking both in the scholarly literature and in the law. This is due in part to the compliance focus in many data protection systems, so that harm was considered collecting personal information without providing proper notice or without obtaining consent, or using data outside of the scope of that consent.

That does not equate with the way most people think about data-related harms, which is more focused on data being used in a way that might cause them injury or embarrassment, rather than the presence or content of privacy notices. So there is a widespread need to think more critically about what constitutes a harm that the risk management framework should seek to minimize or prevent when evaluating data uses.

A framework for recognized harms is critical to ensuring that individuals are protected and enhancing predictability, accountability, and efficiency. National regulators are well placed to help lead a transparent, inclusive process to articulate that framework. The goal should not be to mandate a one-size-fits-all approach to risk analysis, but rather to provide a useful, practical reference point with sufficient clarity to help guide the risk analyses of data user, and to ensure that a wide range of interests and constituencies are involved in crafting it.

Risk assessment is not binary and is likely to be influenced by a number of factors within the data user’s control. So the goal of the risk assessment isn’t simply to indicate whether a proposed data use is likely to be appropriate or not, but also to highlight the steps that the data user can take to make that use more acceptable (*e.g.*, by truncating, encrypting, or de-personalizing data).

E. Transparency and Redress

AI applications are increasingly being used to make decisions about individuals, and even to predict their future behavior, with often significant consequences. Whenever AI is used in ways that affects individuals, there must be effective transparency and redress. This is necessary to protect the rights of individuals, but it also serves the vital purposes of enhancing the accuracy and effectiveness of AI tools, and creating disincentives for deploying tools inappropriately. Meaningful transparency and redress, together with effective enforcement, not only provide remedies for current harms, but also help to prevent future ones.

Moreover, while few individuals demonstrate much interest in inquiring into data processing activities until there is a perceived harm, when they are often more interested in learning how data was used and to what effect. Ensuring that there is meaningful redress will not only create disincentives for risky data processing and help repair the damage that such processing can cause, but it will also provide meaningful rights to individuals at the very time they are most interested in exercising them.

V. Conclusion

The proliferation of AI is already yielding significant benefits, but it also raises important issues. Efforts to address those issues within existing data protection frameworks increasingly demonstrate the limits of those frameworks and their inadequacy both for protecting privacy and for facilitating innovation in an increasingly data-dependent economy. As new AI applications are developed and deployed, we have an opportunity and an increasingly unavoidable need to examine the effectiveness of current data protection laws and to revise them in light of 21st-century realities.

Bibliography

A. Books

Nissenbaum, H., *Privacy in Context* (Stanford University Press 2010).

B. Articles

Camparo, A., This basketball-playing robot is so good it could outshoot Stephen Curry, nbcnews.com (20 Mar. 2018), available at <https://www.nbcnews.com/mach/science/basketball-playing-robot-so-good-it-could-outshoot-stephen-curry-ncna858011>.

Cate, F., "Big Data, Consent, and the Future of Data Protection," in Sugimoto, C., Hamid R. Ekbia & Michael Mattioli, eds., *Big Data is Not a Monolith 2* (MIT 2016).

Cate, F., "Protecting Privacy in Big Data," *Journal of Law and Economic Regulation*, vol. 8, no. 1 (2015).

Chin, M., An AI just beat top lawyers at their own game, Mashable (26 Feb. 2018), available at <https://mashable-com.cdn.ampproject.org/c/s/mashable.com/2018/02/26/ai-beats-humans-at-contracts.amp>.

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/>.

Del Bello, L., AI Translates News Just as Well as a Human Would, futurism.com (16 Mar. 2018), available at <https://futurism.com/ai-translator-microsoft/>.

Devlin, H., "Could online tutors and artificial intelligence be the future of teaching?," *The Guardian* (26 Dec. 2016), available at <https://www.theguardian.com/technology/2016/dec/26/could-online-tutors-and-artificial-intelligence-be-the-future-of-teaching>.

Dwork, C., "Differential Privacy: A Cryptographic Approach to Private Data Analysis," *Privacy, Big Data, and the Public Good: Frameworks for Engagement* at 296 (2014).

Greenberg, A., "An AI That Reads Privacy Policies So That You Don't Have To," *Wired* (9 Feb. 2018), available at <https://www.wired.com/story/polisis-ai-reads-privacy-policies-so-you-dont-have-to/>.

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.

Hiring Based in Neuroscience + Data Science, Pymetrics, available at <https://www.pymetrics.com/science/>.

Hardesty, L., "Is MIT Giving Away the Farm?," *MIT Technology Review* (21 Aug. 2012), available at <https://www.technologyreview.com/s/428698/is-mit-giving-away-the-farm/>.

Keohane, J., "What News-Writing Bots Mean for the Future of Journalism," *Wired* (16 Jan. 2017), available at <https://www.wired.com/2017/02/robots-wrote-this-story/>.

Kirk, J., "EFF: Browsers Can Leave a Unique Trail on the Web," *PC World* (29 Jan. 2010), available at <https://www.pcworld.com/article/188134/article.html>.

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/>.

Kuang, C., "Can A.I. Be Taught to Explain Itself?," *New York Times Magazine* (21 Nov. 2017), available at https://www.nytimes.com/2017/11/21/magazine/can-ai-be-taught-to-explain-itself.html?_r=0.

Kuner, C., et al, "Risk Management in Data Protection," *International Data Privacy Law* 5, no. 2 (2015), 95, available at <https://academic.oup.com/idpl/article/5/2/95/645238>.

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>.

Landau, S., "Control use of data to protect privacy," *Science* 347:6221, Jan. 30, 2015, at 504.

Marr, B., "How Is AI Used In Healthcare—5 Powerful Real-World Examples That Show The Latest Advances," *Forbes* (27 Jun. 2018), available at <https://www.forbes.com/sites/bernardmarr/2018/07/27/how-is-ai-used-in-healthcare-5-powerful-real-world-examples-that-show-the-latest-advances/#5b20b9975dfb>.

Meister, J., AI Plus Human Intelligence Is The Future of Work, *Forbes* (11 Jan. 2018), available at <https://www.forbes.com/sites/jeannemeister/2018/01/11/ai-plus-human-intelligence-is-the-future-of-work/#789369cf2bba>.

Ohm, P., "Changing the Rules: General Principles for Data Use and Analysis," *Privacy, Big Data, and the Public Good: Frameworks for Engagement* at 100 (2014).

Oltsik, J., Artificial intelligence and cybersecurity: The real deal, csoonline.com (25 Jan. 2018), available at <https://www.csoonline.com/article/3250850/security/artificial-intelligence-and-cybersecurity-the-real-deal.html>.

Rieland, R., "Artificial Intelligence Is Now Used to Predict Crime. But Is It Biased?," *Smithsonian* (5 Mar. 2018), available at <https://www.smithsonianmag.com/innovation/artificial-intelligence-is-now-used-predict-crime-is-it-biased-180968337/>.

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/>.

Timmer, J., AI trained to spot heart disease risks using retina scan, arstechnica.com (24 Feb. 2018), available at <https://arstechnica.com/science/2018/02/ai-trained-to-spot-heart-disease-risks-using-retina-scan/>.

Tully, P., Using defensive AI to strip cyberattackers of their advantage, venturebeat.com (6 Mar. 2018), available at <https://venturebeat.com/2018/03/06/using-defensive-ai-to-strip-cyberattackers-of-their-advantage/>.

Vanian, J., "How IBM is Bringing Watson to Wine," *Fortune* (9 Jan. 2016), available at <http://fortune.com/2016/01/09/ibm-bringing-watson-wine/>.

"We'll See You, Anon," *The Economist* (13 Aug. 2015), available at <https://www.economist.com/science-and-technology/2015/08/13/well-see-you-anon>.

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/>.

C. Other References

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 Feb. 2018) at 19, available at http://ec.europa.eu/newsroom/article29/document.cfm?doc_id=49826.

Artificial Intelligence and Privacy, Datatilsynet (Norwegian Data Protection Authority) at 4 (Jan. 2018), available at <https://www.datatilsynet.no/globalassets/global/english/ai-and-privacy.pdf>.

Artificial Intelligence, The Public Policy Opportunity, Intel (18 Oct. 2017), available at <https://blogs.intel.com/policy/files/2017/10/Intel-Artificial-Intelligence-Public-Policy-White-Paper-2017.pdf>.

Big Data, Artificial Intelligence, Machine Learning and Data Protection, United Kingdom Information Commissioner's Office at 11 (Version 2.2 - 2017) (emphasis added), at

<https://ico.org.uk/media/for-organisations/documents/2013559/big-data-ai-ml-and-data-protection.pdf>.

Cate, F., P. Cullen & V. Mayer-Schönberger, *Data Protection for the 21st Century World* (Microsoft 2014).

Cate, F. & V. Mayer-Schönberger, *Data Use and Impact Global Workshop* (Center for Applied Cybersecurity Research 2013).

Cate, F., P. Cullen & V. Mayer-Schönberger, *Data Protection Principles for the 21st Century* (Oxford Internet Institute 2013), available at <https://www.repository.law.indiana.edu/facbooks/23/>.

Centre for Information Policy Leadership, *A Risk-based Approach to Privacy: Improving Effectiveness in Practice* (2014), available at <http://www.hunton.com/files/upload/Post-Paris Risk Paper June 2014.pdf>;

Centre for Information Policy Leadership, *The Role of Risk Management in Data Protection* (2014), available at http://www.hunton.com/files/Uploads/Documents/Centre/Role_of_Risk_Management_in_Data_Protection.pdf.

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.

English Oxford Living Dictionaries, “Artificial Intelligence,” available at https://en.oxforddictionaries.com/definition/artificial_intelligence.

Executive Office of the President of the United States, National Science and Technology Council Committee on Technology, *Preparing for the Future of Artificial Intelligence* (Oct. 2016), available at https://obamawhitehouse.archives.gov/sites/default/files/whitehouse_files/microsites/ostp/NSTC/preparing_for_the_future_of_ai.pdf.

Executive Office of the President, *Big Data: Seizing Opportunities, Preserving Values* xiii (2014), available at https://obamawhitehouse.archives.gov/sites/default/files/docs/big_data_privacy_report_may_1_2014.pdf.

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>.

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/>.

International Organization for Standardization, *ISO 31000:2009 Risk management—Principles and guidelines*.

Machine Learning, Coursera, available at <https://www.coursera.org/learn/machine-learning>.

Organisation for Economic Co-operation and Development, 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.

Organisation for Economic Co-operation and Development, *OECD Guidelines Governing the Protection of Privacy and Transborder Flows of Personal Data*, C(80)58/FINAL, as amended by C92013)79 (2013), 12, available at <http://www.oecd.org/sti/ieconomy/2013-oecd-privacy-guidelines.pdf>.

Organisation for Economic Co-operation and Development, *Supplementary Explanatory Memorandum to the Revised Recommendation of the Council Concerning Guidelines Governing the Protection of Privacy and Transborder Flows of Personal Data* (2013), 30, available at <http://www.oecd.org/sti/ieconomy/2013-oecd-privacy-guidelines.pdf>.

Project InnerEye—Medical Imaging AI to Empower Clinicians, Microsoft Project InnerEye (7 Oct. 2008), available at <https://www.microsoft.com/en-us/research/project/medical-image-analysis/>.

Project Premonition aims to detect pathogens before they cause outbreaks, Microsoft Project Premonition (2 Mar. 2015), available at <https://www.microsoft.com/en-us/research/project/project-premonition/#>.

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 9, available at <https://www.pdpc.gov.sg/-/media/Files/PDPC/PDF-Files/Resource-for-Organisation/AI/Discussion-Paper-on-AI-and-PD---050618.pdf>.

South Korea Personal Information Protection Act, Article 2(1). 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#>.

Spencer, G., Buffaloes and the Cloud: Students turn to tech to save poor farming families, news.microsoft.com (27 Sep. 2017), available at <https://news.microsoft.com/apac/features/saving-farming-families-tech-one-cow-goat-buffalo-time/>.

Stone, P., et al, "Artificial Intelligence and Life in 2030." *One Hundred Year Study on Artificial Intelligence: Report of the 2015-2016 Study Panel*, Stanford University (Sep. 2016)
https://ai100.stanford.edu/sites/default/files/ai100report10032016fnl_singles.pdf.

Sweeney, L., Simple Demographics Often Identify People Uniquely, Carnegie Mellon University, Data Privacy Working Paper at 3 (2000), available at
<https://dataprivacylab.org/projects/identifiability/paper1.pdf>.