Intel's AI Privacy Policy White Paper

Protecting individuals' privacy and data in the artificial intelligence world

Authors:

David Hoffman
Associate General Counsel and Global Privacy Officer

Riccardo Masucci
Global Director of Privacy Policy

Executive Summary

Two Technology Trends

At Intel, we believe in the enormous potential of artificial intelligence (AI) to improve people's lives. Two trends that are influencing AI and digital developments globally are: **1. Data analytics from the edge to the cloud**, and **2. Increased mechanisms for data collection and creation**.

Five Foundational Observations on AI and Privacy

AI represents a new privacy territory as it entails autonomous determinations that potentially affect citizens. The following observations shape privacy's impact on AI:

a) Increased automation should not translate to less privacy protection;

b) Explainability needs more accountability;

c) Ethical data processing is built on privacy;

d) Privacy protects who we are (how others see us and how we see ourselves);

e) Encryption and de-identification help address privacy in AI.

Six Policy Recommendations

1	**New legislative and regulatory initiatives should be comprehensive, technology neutral, and support the free flow of data:** horizontal legislation can encompass both data uses and technologies that fall outside existing sectoral laws and that are still unforeseen.
2	**Organizations should embrace risk-based accountability approaches**, putting in place technical (privacy-by-design) or organizational measures (product development lifecycles and ethics review boards) to minimize privacy risks in AI.
3	**Automated decision making should be fostered while augmenting with safeguards to protect individuals:** legitimate interest should be acknowledged as legal basis for data processing for AI. Industry and governments should work together on algorithm explainability and risk-based degrees of human oversight to minimize potential adverse impacts for citizens from automated decision-making.
4	**Governments should promote access to data**, for example, opening up government data, supporting the creation of reliable datasets available to all, fostering incentives for data sharing, investing in the development of voluntary international standards (i.e. for algorithmic explainability) and promoting diversity in datasets.
5	**Funding research in security is essential to protect privacy:** in areas like homomorphic encryption, access to personal data can be minimised and protection enhanced.
6	**It takes data to protect data:** to detect biases or cyber threats and to protect personal data, AI needs to process personal data.

I. Two Technology Trends

Artificial intelligence (AI) has gained prominence in the public debate due to the tremendous potential of its applications.[1] Innovation across the digital society we live in is astounding: autonomous technologies are deployed for new life-enhancing and potentially life-saving uses such as disease detection, precision medicine, driving assistance, increased productivity, safety at work and to make education more accessible.

These advancements can be achieved thanks to increased computing capability that supports complex algorithms to extract meaningful information from ever-bigger datasets. Compute power and data are the enablers of artificial intelligence. Data represents the fabric of our contemporary world and its quantity has grown exponentially over the past few years. Some estimated that every day 2.5 quintillion bytes of new data are produced, and more surprisingly, 90% of all available data was produced after 2015.[2]

In such a data-intensive society, two trends will affect future developments:

1) **Edge to cloud data analytics:** Essential building blocks of the environment described above are data centres and networks that serve millions of connected "edge devices", from smart homes and factories, to autonomous cars and drones. These endpoints leverage increasing capabilities in connectivity (such as future 5G communications) and computing power to carry out analytical workloads to extrapolate value from datasets. Data processing and analytics happen across the infrastructure, at the edge, on the network and in the data centre.
2) **Increased mechanisms for data collection and creation:** Personal data is not just collected from individuals who provide it for particular uses, but also observed and gathered by sensors in connected devices, and derived or created through further automated processing.[3] In fact, the percentage of data coming directly from individuals is decreasing compared to the information that is collected in our increasingly connected society and inferred through machine learning technologies.[4]

The unprecedented opportunities linked to development and adoption of AI-based solutions are drawing policymakers' and regulators' attention to implications for citizens and society, as well as to possible actions that governments and national authorities can take. Privacy and data protection represent a key component of these implications and possible government actions. The interest in AI is growing across the world and many major economies have defined or are in the process of shaping their national AI strategies.

II. Five Foundational Observations on AI and Privacy

Due in large part to those two trends, increasingly autonomous and ubiquitous technologies take advantage of large datasets and data from multiple sources to make autonomous determinations in near-real time. In some cases (i.e. banking, human resources, transportation), these decisions may affect individuals, their private lives, their physical safety, their position in society, and their interaction with others. It is important for governments to analyze the impact of this automated decision making on privacy.

The existence of potential harm for individuals resulting from autonomous determinations (e.g. discrimination and restriction of choices and possibilities) will create a number of unique situations that public and private organisations will have to deal with when shaping their privacy policies and strategies. Below we describe five observations which inform and inspire our work on AI and Privacy.

a. Increased automation should not translate to less privacy protection.

The OECD Fair Information Practice Principles (FIPPs) have inspired lawmaking in the field of privacy and data protection for the past forty years.[5] Intel's call to policymakers and industry to "Rethink Privacy"[6] is valid now more than ever due to the increasing pressure to which those principles ("the global common language of data protection") are subject. The FIPPs have managed to adapt to technology changes and still

provide valuable guidance because they reflect long-held, widely-accepted values about the individual's relationship with personal data and organizations' responsibility to protect that data. At Intel we believe that the FIPPs should be seen as a system of levers to be pulled and adjusted to provide the best protection possible in the context of a particular data application.[7] When certain FIPPs are less helpful, such as Purpose Specification and Collection Limitation, more investment may be necessary in other FIPPs. With regard to AI, heightened focus should be placed on the FIPPs of Security Safeguards and Accountability.

b. Explainability needs more accountability.

Automated determinations will have an impact on people's lives and their possibility of self-determination. Deep learning techniques (broadly deployed today for applications like computer vision, natural language processing or facial recognition) use large datasets to iteratively train multi-layered neural networks - a process inspired by the human brain. The complexity and abstraction of these tasks lead to situations where the factors driving the results of the algorithms are hard to understand and therefore explain. The principle of transparency will be challenged because the logic involved in the decision making will be difficult to discern. While some academics assume the implicit existence of a "right to explanation" in the EU GDPR[8], organisations implementing AI solutions should be able to demonstrate that they have the right processes, policies and resources in place to minimize privacy risks and adverse impacts to the individual.

Responsible risk management practices require that organisations hold themselves accountable to put in place appropriate technical and organisational measures for addressing privacy and data protection concerns of customers, business partners and society. Accountability starts with the organisation's commitment to create tools and training to implement privacy, to promote systems for internal oversight and external verification, and to ensure means for remediation and enforcement.[9] Privacy-by-design approaches, that embed privacy impact assessments throughout the design and production process, are a good example of accountability and should apply across the AI ecosystem.

c. Ethical data processing is built on privacy.

Intel has long promoted the innovative and ethical data use for the transformative positive impact it can make on the lives of individuals. The continued evolution of our digital society coincides with a series of socio-cultural shifts, such as from human autonomy to the convergence between human and machines. Similarly, traditional values such as dignity, freedom, democracy, equality, autonomy, and justice are part of discussions around digital ethics.[10] Privacy and data protection today represent intrinsic and foundational concepts for our modern society, which enable individual freedom of choice and user control. Protecting individuals and their data goes beyond legal compliance requirements: it means embracing societal values and working to build a much-needed trust in the technologies and their positive impact on people.

d. Privacy is protecting who we are.

Privacy requires that data is both reliable and that it will not be used to harm individuals. Privacy aims to prevent unauthorised access, modification and loss of personal data. It requires a respect for private and family life, home and confidentiality of communications. Artificial intelligence techniques and their potential to create data such as images, videos, and sounds are moving the debate away from the pure risk of identification or unauthorized processing of personal identifiable information. Artificial intelligence is increasing the risks to individuals in two specific ways that impact 1. how others see us, and 2. how we see ourselves.

How others see us: artificial intelligence may manipulate individuals and their reality, creating false information, such as false physical or psychological projections of people. Safeguarding truthfulness of personal data entails the right of citizens to control the use of their name, image or aspect of their identity (personality rights). AI could increase the ability to create false information about individuals that is increasingly believable by others and can impact the way a person is perceived and treated. For these reasons, maintaining the authenticity of personhood

represents a key privacy priority for our society.

How we see ourselves: artificial intelligence techniques boost the ability to perform accurate profiling of people and to extract meaningful insights about them, with unprecedented degrees of depth. Organisations may end up knowing or predicting individuals' behaviours better than people know themselves. Some ill-intentioned organizations and governments could take advantage of this ability, targeting citizens and change their view of reality. This could lead to a "programmable citizen" scenario, where inappropriate use of AI can influence individuals' choices (for example, for electoral purposes through fake news) or modify what groups they identify with (for example, for control and surveillance purposes in non-democratic regimes).

e. Stronger encryption and de-identification help address privacy.

Data confidentiality, integrity and accuracy have been key objectives of cybersecurity practices for decades and have been contemplated in all privacy legal frameworks. With the increasing amount of data collected, processed and inferred in the artificial intelligence space, strong encryption and de-identification (full anonymization) techniques serve the purpose of protecting individuals' privacy while achieving higher levels of security. Achieving de-identification will require increasingly complex practices, because re-identification will be increasingly possible in a deep learning-driven environment. Differential privacy techniques have emerged in the last years as viable solutions to minimise privacy risks, adding "noise" to scramble personal data.[11] In the academic and research community, homomorphic encryption seems particularly promising as it allows computation on encrypted data, therefore enabling AI tasks without the need to transfer personal information.[12]

III. Six Policy Recommendations

This paper aims at informing and contributing to Intel's ongoing engagement with policymakers and regulators to find effective solutions to address privacy concerns. Intel works at developing "edge-to-cloud" hardware and software solutions that will enable artificial intelligence to realize its potential for improving the lives of people. At the same time, Intel is committed to allow citizens to benefit from these innovative data uses, while trusting that their data is processed in ethical, protected and privacy-preserving ways.

1. New legislative and regulatory initiatives should be comprehensive, technology neutral and should enable the free flow of data.

To address the complexity of our digital society and keep up with the pace of technology advances, privacy legislation should be comprehensive to avoid legal loopholes and to cover both data uses and technologies that currently either fall outside existing sectoral laws or that are still unforeseen. AI-specific privacy laws might not ensure enough flexibility to pass the test of time. Instead, comprehensive, horizontal legislation should be based on the FIPPs with flexibility provided to increase protections from certain FIPPs when others are not sufficiently protective of individuals' privacy. FIPPs such as Purpose Specification, Collection Limitation, and Openness will be challenged by autonomous technologies; therefore, any proposed law should focus on effective solutions allowing for security safeguards, individual participation, and increased accountability of organizations. The history of the FIPPs - which has inspired privacy legislation around the world - shows how long standing principles have been and can be used and reinterpreted according to technological advancements to create technology neutral laws.

In addition, the ability to process data and move it across borders is critical to developing new technologies. The global value of digital data flows has grown exponentially over the past decade and accounts for several trillions of dollars of the global GDP.[13] Legislative initiatives should, therefore, promote unencumbered cross border data flows under appropriate safeguards

for individuals. At the same time, we acknowledge that governments may have a number of legitimate reasons to require the storing of a copy of the data within their country. Defining those situations, would not prevent policymakers from supporting the protected free flow of data and limit data localisation requirements.

2. Organizations should embrace risk-based accountability approaches

Accountability can be described as the ability of responsible organisations to demonstrate that appropriate measures have been put in place to minimize privacy and security risks. These technical or organizational measures should be tailored based on each business' needs as well as the specific risks associated to the data processing performed. Consequently, regulators could deem accountability measures as a proof of legal compliance or at least a mitigating factor in case of a data breach.

From a technical perspective, embedding privacy in the development of AI solutions (privacy-by-design) is possible. For example, anonymous video analytics software can extract information, store it as a log file and discard the images. Image blurring of faces or car plates minimize the risk of linking identifiers to individuals. Development of voluntary international standards to guide common privacy-by-design approaches could enable consistency across multiple jurisdictions. In addition, de-identification and encryption practices have proved effective in addressing privacy and security concerns, because they still enable computational tasks without exposing data to unauthorised access, modification, loss or deletion. These practices are crucial in the health care sector where medical information represents sensitive data that, if misused, could harm patients or subject them to discrimination.

Responsible companies have developed product development lifecycles, which include impact assessments and balancing tests to measure privacy and security risks of artificial intelligence-based technologies. These assessments are the means by which accountable organizations can deliver on privacy-by-design. Throughout the production process, from design to market release, engineers and analysts evaluate at each step potential unintended consequences for end users, and suggest additional features or alternative organisational solutions to minimise those risks. For example, researchers may adapt their data management practices, until they understand the implications of that specific use of the data; they may segment or cut out information before passing it on to other data scientists; they may store data in few highly secured locations with access restricted to a small group of users.

Accountable organizations also need governance structures over these design processes. Privacy-conscious and ethical use of data has become a priority also for senior management in organisations. Future technology scenarios are increasingly discussed within ethics councils or ethical review boards established by organizations to address ethical issues inherent to products that are essential to build and reinforce citizens' trust..

Intel has already taken significant steps towards all the above-mentioned measures and firmly believes that industry across the board should embrace similar approaches.

3. Automated decision-making should be fostered while augmenting with safeguards to protect individuals.

The "notice and consent" model has attempted to provide for the FIPP of Individual Participation for decades. It has been a valuable tool to empower citizens and give them control over data, but has always had limited effect due to the tremendous burden it places on individuals to fully understand how information that relates to them is collected, processed and used. The ever-changing technology environment shows that the use of notice and consent will increasingly be difficult in many data collection and creation contexts, hence other legal bases should be considered as lawful grounds for processing. The legitimate interest of the entity processing personal data (controller or processor) should be balanced against the legitimate expectations of the individuals (data subjects). Well-identified legitimate interests

such as the physical safety of individuals or network and information security should supplement consent where appropriate. Legitimate interest can serve as an effective legal basis for data processing when genuine, explicit consent cannot be obtained. It does not represent a blank authorization to process data, but it works in concert with the other substantive rights provided to individuals (access, correction, deletion, portability) and obligations on organizations, such as security safeguards and accountability approaches (described in the following subsection)

Automated decision-making should not be limited a priori because this approach risks impinging on innovation and possibly preventing citizens from life enhancing AI applications. For example, a restrictive interpretation of article 22 of the EU GDPR (which prohibits determinations based solely on automated decision-making) might negatively affect basic functions and future developments of autonomous technologies, such as autonomous driving. In this example, legal basis exceptions like consent or contracts might not cover the whole spectrum of cases where, for example, passengers, other drivers, and/or pedestrians are involved.

Different levels of human involvement (human oversight) could be tailored to different autonomous functions. AI systems should be designed, built, and deployed to allow for the control and judgment of humans for those AI applications that present higher risk profiles for individuals impacted by decisions. We believe that risk-based degrees of human oversight could represent a practical solution to allow for innovative uses of AI and to ensure control over autonomous decision-making when necessary. AI-based industrial quality control mechanisms may not deserve human oversight when they single out a defective product. On the contrary, if the algorithm excludes a candidate in a job selection procedure or a patient for treatment eligibility, there would be an expectation from individuals and a public interest to have at some point a human involvement in or oversight of the decision-making.[14] Accountable organisations would proactively seek solutions to ensure the right degree of human intervention is achieved based on the potential risks for individuals.

4. Governments should promote access to data.

Access to large and reliable datasets is essential to the development and deployment of AI. Improving access to reliable data would be beneficial for making design more competitive and innovative, for better achieving public policy priorities and for increasing quality and quantity of products and services available to citizens. The current situation could be improved by government initiatives aiming at:

a) Making available public sources of information in structured and accessible databases (open government data).

b) Actively supporting the creation of reliable datasets (including personal information of individuals), which could be used by all AI developers, by start-ups and more broadly by industry to test automated solutions and benchmark the quality of their algorithms.

c) Fostering incentives for data sharing between the public and private sector and among industry players. Similar to the efforts made in the cybersecurity space for threat information, appropriately shared and protected data would increase awareness and improve privacy of the whole automated environment. A good example could be governments funding the creation of APIs (Application Programming Interfaces) to allow for dynamic access to data.

d) Contributing to the creation of international voluntary standards that allow for easier information sharing while noting which data fields could include personal information (so that they can be reviewed for privacy issues); provide guidance on methods for responsible data sharing; and define algorithm explainability for different AI implementations.

e) Promoting diversity in datasets. The potential presence of biases in the design of algorithms feed users' distrust in technology. Greater diversity will reduce the risk of unintended bias. In AI for medicine, it is important that data from ethnic minorities is included in the training data, so the result is more accurate and doesn't preclude appropriate medical treatment to portions of society. The ability for research organisations to transfer data internationally reinforces this opportunity and governments should encourage cross border data flows.

5. Funding research in security is essential to protect privacy.

The pace of AI advancements in the past few years has been unprecedented. However, more work still needs to be done to improve computing power, energy efficiency and connectivity in data centers and in the edge devices. More research is needed in areas such as homomorphic encryption, which will allow extracting meaningful information from encrypted data without the need to disclose personal information and, therefore, will protect individuals' privacy.

6. It takes data to protect data.

The tremendous potential of AI includes also the possibility to deploy analytics to support and serve public policy priorities such as privacy, data protection and cybersecurity. In fact, algorithms can help detect unintended discrimination and bias, identity theft risks or cyber threats. Artificial intelligence provides organisations the ability to prevent and manage risks, including threats of large scale privacy breaches. In this way, the use of this data (some of which will be personal data) actually has a net positive impact on privacy.[15]

IV. Conclusion

If the benefits of artificial intelligence and its numerous applications seem indisputable, further debate is actually needed to address potential unintended societal consequences and risks for individuals related to the use of AI. This paper outlines how some unique characteristics of AI may lead to specific privacy risks. Acknowledging these concerns, we believe that the five foundational observations described in this paper will help realize the full potential of AI to benefit people. The list of six policy recommendations is aimed at starting a fruitful discussion with policymakers, regulators and like-minded organisations that are exploring public policy solutions to data protection and privacy challenges. We welcome the opportunity to discuss these issues further as well as to receive feedback on our proposals. Intel stands ready to work closely with all interested stakeholders to develop viable pathways to pursue AI innovation and privacy together.

References

1. "AI is best understood as a set of techniques aimed at approximating some aspect of human or animal cognition using machines." Ryan Calo, Artificial Intelligence Policy: A Primer and a Roadmap, 51 UC Davis Law Rev. 399, 404 (2017).
2. Source IBM marketing https://public.dhe.ibm.com/common/ssi/ecm/wr/en/wrl12345usen/watson-customer-engagement-watson-marketing-wr-other-papers-and-reports-wrl12345usen-20170719.pdf
3. A useful taxonomy on data origins was developed by Marty Abrams, Information Accountability Foundation (IAF), in 2013 http://informationaccountability.org/wp-content/uploads/Data-Origins-Abrams.pdf
4. Some data suggesting this trend can be found in the IDC report, "Data Age 2025: The evolution of Data to Life-Critical", 2017
5. The eight FIPPs are Collection Limitation, Data Quality, Purpose Specification, Use Limitation, Security Safeguards, Openness, Individual Participation and Accountability. OECD Revised Guidelines on the Protection of Privacy and Transborder Flows of Personal Data (2013), http://oecd.org/sti/ieconomy/oecd_privacy_framework.pdf
6. D. Hoffman, P. Bruening, 2014 Rethink Privacy: fair information practices reinterpreted https://blogs.intel.com/policy/files/2015/01/RethinkingPrivacy.pdf
7. P. J. Bruening and H. M. Patterson. A Context-Driven Rethink of the Fair Information Practice Principles. SSRN Artificial Intelligence — Law, Policy, and Ethics eJournal. Available at http://ssrn.com/abstract=2843315
8. Selbst, Andrew D. and Powles, Julia, Meaningful Information and the Right to Explanation. International Data Privacy Law, vol. 7(4), 233-242 (2017). https://ssrn.com/abstract=3039125
9. Centre for Information Policy Leadership, Data Protection Accountability: The Essential Element A Document for Discussion, 2009 https://www.informationpolicycentre.com/uploads/5/7/1/0/57104281/data_protection_accountability-the_essential_elements__discussion_document_october_2009_.pdf
10. See European Data Protection Supervisor (EDPS) Ethics Advisory Group's report on this topic: https://edps.europa.eu/sites/edp/files/publication/18-01-25_eag_report_en.pdf
11. Differential Privacy is a system for limiting the privacy loss that an individual experiences when their private information is used to create an aggregate data product. Differential privacy mechanisms operate by introducing randomness into the results of queries on underlying confidential data. Because of the randomness, an observer of the queries faces ambiguity when trying to reconstruct what the confidential data must have been in order to produce the observed results. (Source: Wikipedia)
12. Homomorphic Encryption is a form of encryption that allows computations to be carried out on ciphertext (encrypted information), thus generating an encrypted result which, when decrypted, matches the result of operations performed on the plaintext. The purpose of homomorphic encryption is to allow computation on encrypted data. Cloud computing platforms can perform difficult computations on homomorphically encrypted data without ever having access to the unencrypted data. (Source: Wikipedia)
13. McKinsey Global Institute (MGI) report, "Digital globalization: The new era of global flows", February 2016 https://www.mckinsey.com/~/media/McKinsey/Business%20Functions/McKinsey%20Digital/Our%20Insights/Digital%20globalization%20The%20new%20era%20of%20global%20flows/MGI-Digital-globalization-Full-report.ashx
14. An interesting perspective was offered by the Singapore Data Protection Commission in the "Discussion Paper on Artificial Intelligence", which foresees different degrees of involvement for human decision makers, depending on the severity and probability of harm to individuals. Humans can be *in-the-loop* (making the final decision), *over-the-loop* (taking the decision based on options defined by AI), or *out-of-the-loop* (excluded, but in a predefined set of use scenarios). https://www.pdpc.gov.sg/-/media/Files/PDPC/PDF-Files/Resource-for-Organisation/AI/Discussion-Paper-on-AI-and-PD---050618.pdf
15. D. Hoffman, P. Rimo, "It takes data to protect data", 2018 https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2973280

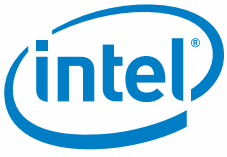

22 October 2018