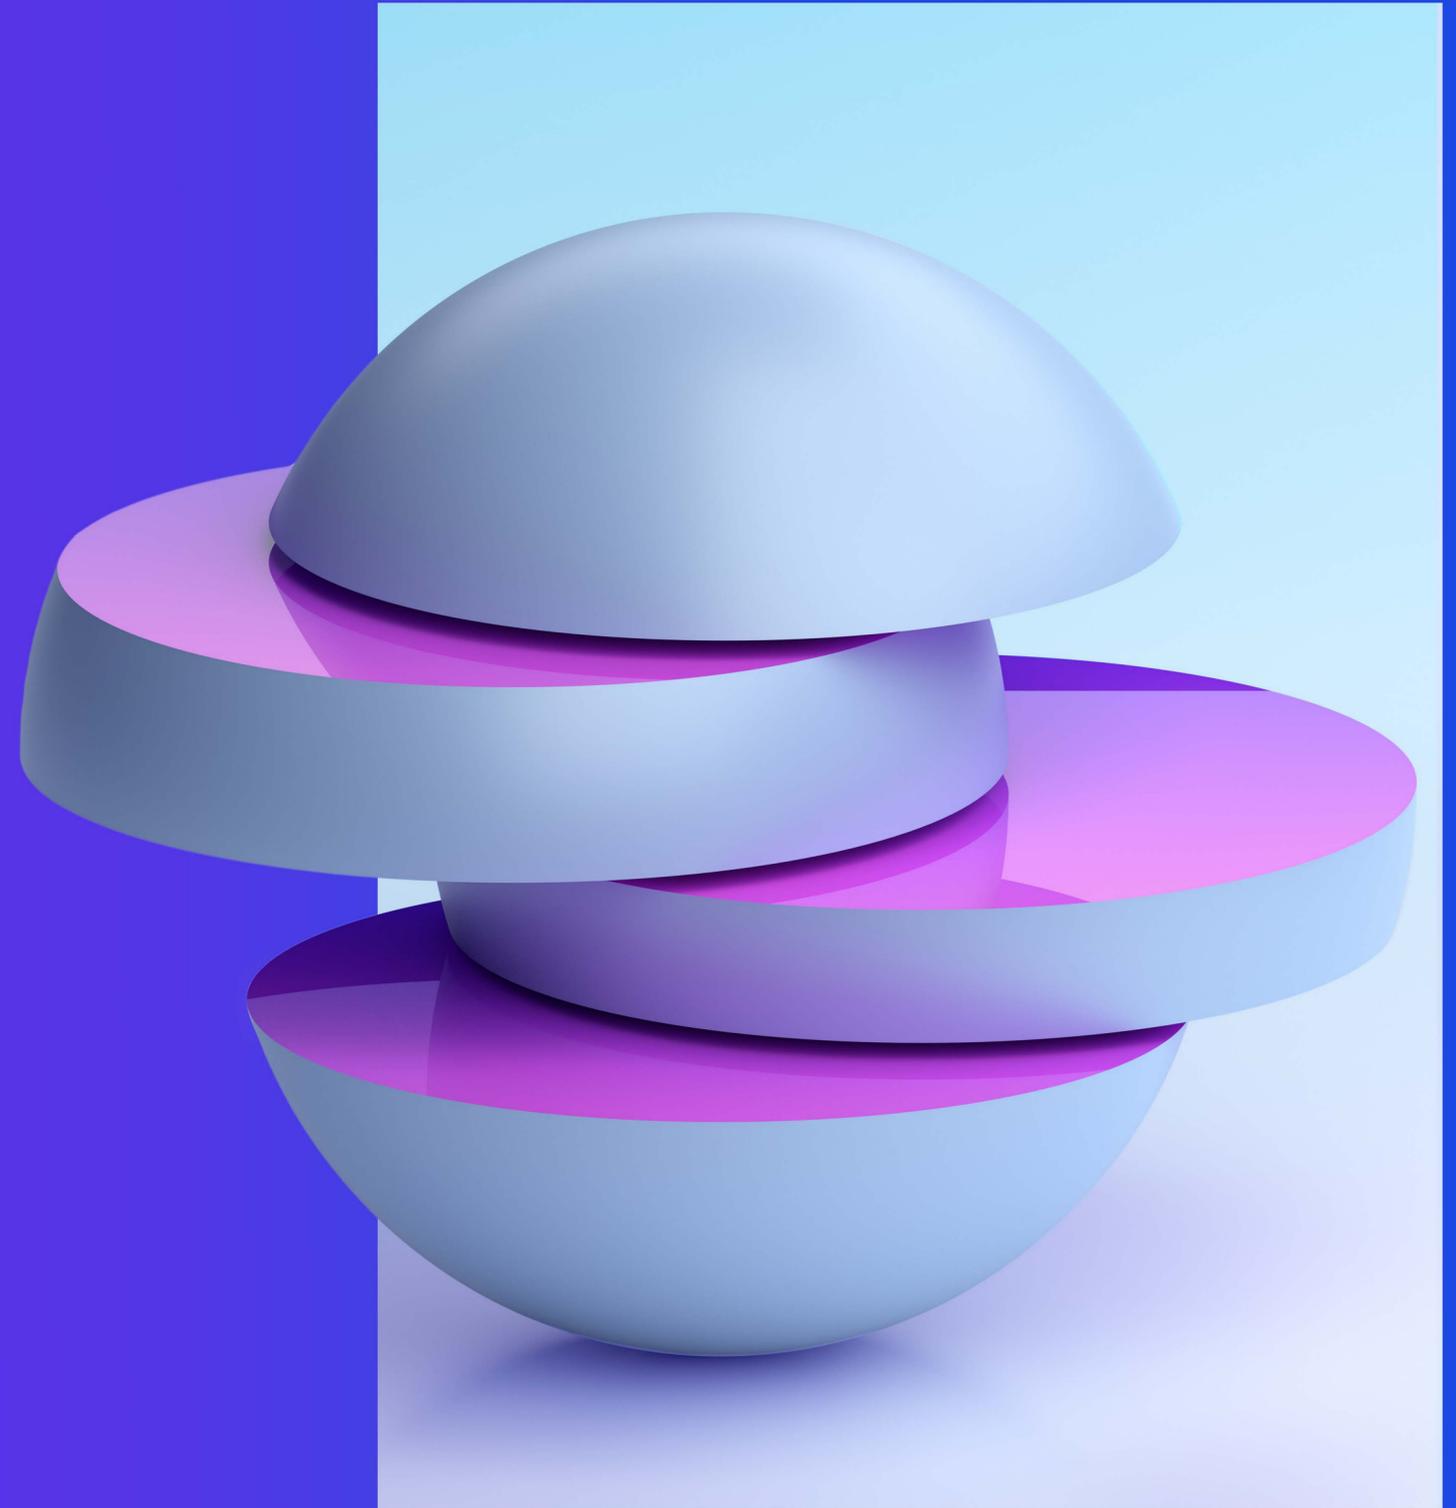




Privacy in the new world of AI

How to build trust in AI through privacy.





Contents

03

Data privacy and trust in AI — the promise

04

What's next? The evolving regulatory landscape

06

Key principles for achieving AI privacy

11

Regulatory highlights

12

The road ahead: Building trustworthy AI





Data privacy and trust in AI — the promise

Artificial Intelligence (AI) promises to transform our lives, helping to make us more efficient, productive, healthier and innovative.

This exciting technology is already being used across the private and public sectors, harnessing the power of data to improve forecasting, make better products and services, reduce costs, and free workers from routine administrative work.

In the healthcare sector, for example, doctors can more accurately and rapidly predict health risks and carry out complex treatments more effectively. In mining, AI-powered robots are undertaking hazardous tasks such as coal mining, sea exploration, and also helping in rescue operations during natural disasters. In commercial banking, AI and ML are helping sales

and marketing teams identify prospects and predict customer needs and buying propensity. They also enable dynamic deal pricing for micro-segments as well as automate decision-making processes, credit rule sets and exceptions. In consumer and retail, AI is helping predict and analyze trends, create virtual models that can display outfits, anticipate customer needs and help customers enjoy a more personalized shopping experience.

According to the [KPMG Global Tech Report 2023](#), technology leaders identify AI and machine learning as the most important technology for achieving short-term ambitions.¹ In addition, in the

[2023 Trust in artificial intelligence](#) global survey of more than 17,000 people worldwide, 85 percent believe AI can bring a range of benefits.²

However, as with any emerging technology, there are risks. The same survey finds that 61 percent of people are wary about trusting AI systems, and only half believe the benefits of AI outweigh the risks.³ Also, 55 percent of technology leaders say their progress with automation is delayed because of concerns about how AI systems make decisions.⁴ The widespread and unregulated use of this technology raises concerns about its impact on human rights and personal

privacy. This is especially true for generative AI (GenAI), which is driven by algorithms based on large language learning models (LLMs).

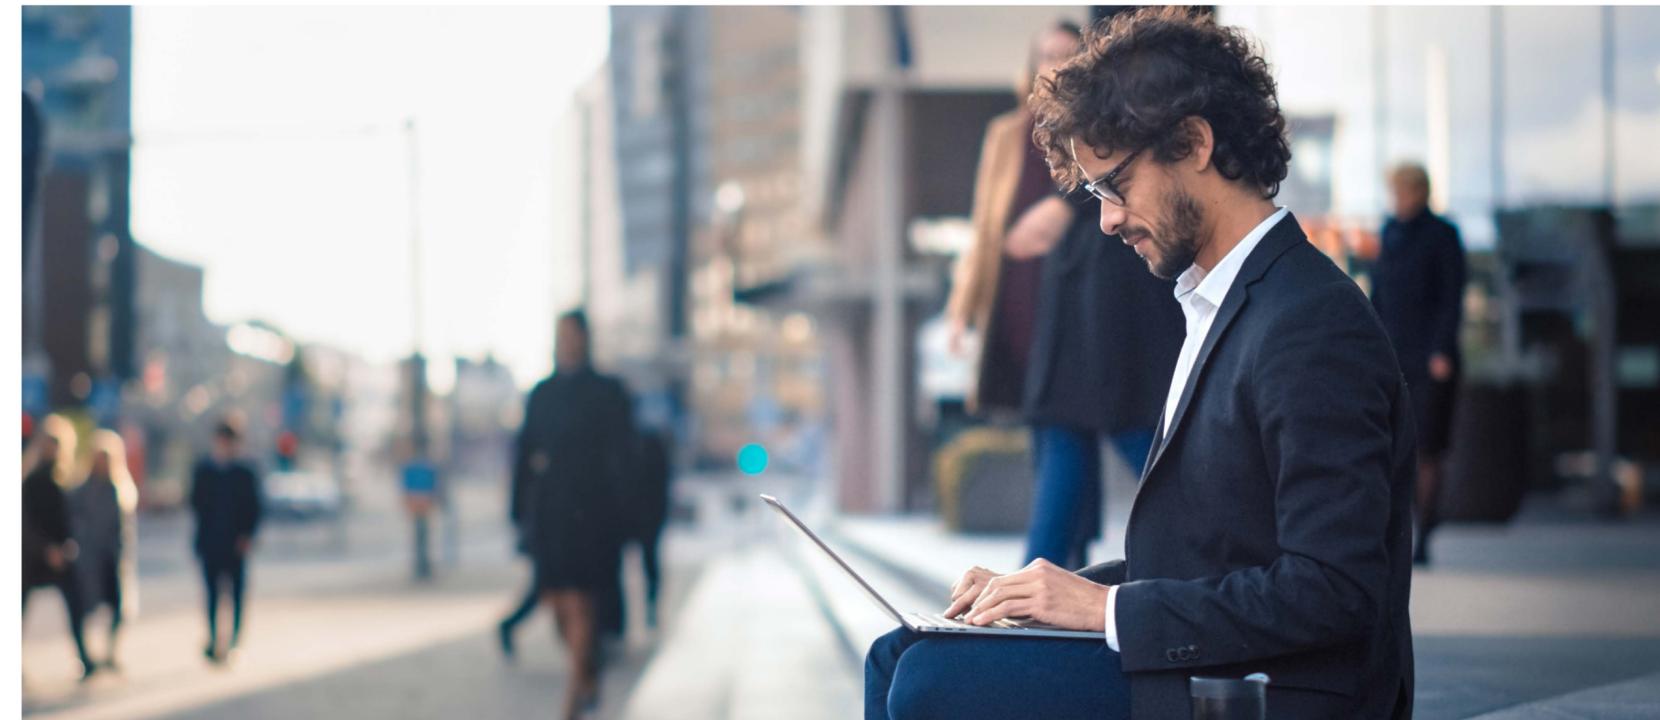
At the time of writing, AI leaders have issued open letters seeking a pause on development in GenAI, urging legislators to future-proof its use through guardrails.⁵ Some cited risks include flawed design, biased logic, coding errors, security vulnerabilities and, most importantly, judgments that discriminate against individuals or groups (after all, the data used was originally created by humans and can reflect existing biases in society). In addition, AI models can generate inaccurate

outputs leading to fake news or disinformation.

Algorithms are also unpredictable, complex and difficult to explain. Given its proprietary nature, they lack transparency and generate outputs based on large-scale data processing from the internet, amplifying the risk of confidential data leaks and breaching legally protected personal information. International privacy laws apply to the collection of data throughout every stage of an AI's lifecycle, so it's little surprise that AI data mining and harvesting has attracted global regulatory scrutiny, with the European Data Protection Authorities and privacy

watchdogs around the world launching investigations into the lawfulness of data processing activities related to GenAI.

This paper investigates the privacy implications of the widespread adoption of AI. It aims to uncover what this means for businesses from a privacy perspective and outlines the key steps organizations can take to utilize AI responsibly. By staying informed about the privacy implications of AI adoption and taking proactive steps to mitigate risks, companies can harness this technology's power while safeguarding individuals' privacy.



¹ KPMG Global Tech Report 2023, KPMG International, 2023.

² Trust in artificial intelligence, The University of Queensland and KPMG Australia, 2023.

³ Ibid.

⁴ KPMG Global Tech Report 2023, KPMG International, 2023.

⁵ Pause Giant AI Experiments: An Open Letter, Future of Life Institute, March 22, 2023.



What's next? The evolving regulatory landscape

As the use of GenAI grows, nations are rushing to legislate and create standards for the responsible use of AI.

The European Union is making a first-world attempt at legislating what could end up being the toughest AI privacy legislation globally. Moreover, the European Data Protection Board (EDPB) has already established a dedicated ChatGPT task force to foster cooperation and information exchange on possible enforcement

activities by data protection authorities. At the time of writing, the European Council, Parliament and European Union (EU) Commission have proposed the Artificial Intelligence Act (AI Act), which, when enacted, could become the toughest global AI privacy legislation.

⁶ Open the black box of algorithms: A deep dive into ECAT's work, European Centre for Algorithmic Transparency, European Commission, April 18, 2023.

The AI Act builds on the privacy provisions under the General Data Protection Regulation (GDPR) that include principles around openness, fairness, algorithmic decision-making and human dignity. These principles formed the basis of the Organization for Economic Co-operation and Development (OECD) principles created in 2019 for responsible stewardship of trustworthy AI, which state that AI systems should be robust, secure and safe throughout its lifecycle.

The AI Act further strengthens the OECD principles by mandating that AI be legally, ethically and technically robust while also respecting democratic values, human rights and the rule of law. Accuracy, non-discrimination, human oversight and attestation to AI standards are all required. The AI Act introduces risk categories and sets legal requirements for 'high-risk' AI systems, including those used in performance evaluations, recruiting and promotions based on the level of public risk posed.

The Digital Services Act (DSA) has also been passed by the EU, which will go into full effect in 2024. The DSA imposes transparency, risk assessment and algorithmic accountability obligations on AI platforms subject to rigorous transparency audits. The European Centre for Algorithmic Transparency (ECAT) will help enforce the DSA and aims to be a critical player in a new digital enforcement



ecosystem within the EU that will have global implications, setting a moral standard for the world.⁶ While regulatory developments for AI are also taking shape in Canada, China and Brazil at the federal level, a single federal approach to legislate AI in the US has not yet emerged despite 131 AI bills having been proposed by the US Congress. However, local state-level regulations for independent

audits, model fairness and transparency for using algorithmic systems have already been enacted in certain states, including Illinois, Maryland, Washington, New York and California. The regulatory regimes currently being established to legislate AI significantly overlap with existing regulatory privacy regimes. These regimes are driven by global concerns about the impact of AI

on individuals, particularly regarding fairness, explainability, transparency, security, respect for the person and accountability. Many of the OECD AI principles developed in 2019 can be mapped to privacy principles in the context of personal information protection and extended to Privacy by Design principles when architecting AI systems.



How the OECD AI principles map to the GDPR privacy principles and the Privacy by Design principles.

OECD AI principles					
	• Accountability	• Inclusive growth, sustainable development and wellbeing	• Human-centered values and fairness	• Transparency and explainability	• Robustness, security and safety
GDPR principles	<ul style="list-style-type: none"> • Accountability 	<ul style="list-style-type: none"> • Purpose limitation • Data minimization 	<ul style="list-style-type: none"> • Lawfulness, fairness and transparency • Purpose limitation • Accuracy • Storage Limitation 	<ul style="list-style-type: none"> • Lawfulness, fairness and transparency 	<ul style="list-style-type: none"> • Integrity and confidentiality
Privacy by Design principles	<ul style="list-style-type: none"> • Proactive not reactive; preventative not remedial 	<ul style="list-style-type: none"> • Privacy embedded into design • Full functionality — positive-sum, not zero-sum 	<ul style="list-style-type: none"> • Privacy as the default setting • Respect for user privacy — keep it user-centric 	<ul style="list-style-type: none"> • Visibility and transparency — keep it open 	<ul style="list-style-type: none"> • End-to-end security — full lifecycle protection

The future of AI regulations

The concurrent development of global AI standards and responsible AI frameworks will shape the future of AI regulations. These standards and frameworks can help provide organizations with a more holistic approach to AI and privacy risk management — especially jurisdictions late to the legislative game.

For example, the US National Institute of Standards and Technology (NIST) has introduced a voluntary, self-regulated Risk Management Framework that aims to ensure the trustworthiness of AI while mitigating risk. The International Standards Organization (ISO/IEC) has also developed ISO 42001 on Artificial Intelligence Management System (AIMS).

When combined with relevant legal rules for AI systems, these standards can serve as effective risk management tools to help operationalize and promote responsible AI best practices and AI governance approaches. They will encourage consistency across different regulatory regimes, which may have conflicting requirements regarding AI systems' robustness, safety, fairness and trustworthiness.

The current US Administration has secured commitments from major technology companies to follow many of the principles and requirements set out in these standards. Time will tell how this will play out in tandem with global regulations.



Key principles for achieving AI privacy

Embedding Privacy by Design into AI systems should help to build trust and navigate potential privacy challenges.

Trust is a key enabler of revenue and growth. Organizations that utilize AI should embed privacy into AI development processes and AI systems to ensure that they are safe, effective and unbiased, supported by strong governance, clear accountability and robust oversight. While waiting for legislation to catch up to technology, organizations wanting to launch AI should embed privacy into every stage of the AI lifecycle as a best practice. Adopting a Privacy by Design approach can help assure customers, regulators and other stakeholders of AI's trustworthiness and minimize any negative impacts.

To help organizations take a proactive Privacy by Design engineering approach, here are key privacy principles to consider throughout the AI lifecycle.

01

Lawfulness and fairness

The AI has a legitimate, legal and clearly-defined purpose with minimal impact on privacy.

Privacy risk

Model design and implementation failures and missing safety features.

These risks can occur when the design specifications are unsuitable for the intended tasks, possibly due to poor choices for inputs and target variables, faulty assumptions or prioritizing efficiency over efficacy.

What could go wrong

A public sector agency established an AI-driven system for detecting benefits fraud that delivered inaccurate, random and unfair results. Amongst the fundamental flaws in model design and implementation was the requirement for invasive, irrelevant, banal and subjective inputs. For example, it used relationship status, duration and frequency to attach high-risk scores for certain applicants but not others. The model also made unjustifiable correlations between different inputs that undermined the consistency of the decisions. Finally, the decision-making process lacked a functional appeals process (a key safety feature).

As a result, the algorithm disproportionately assigned high fraud-risk scores to the people who needed benefits the most without any evidence or compelling justification. Qualified applicants who were flagged were subjected to invasive and stigmatizing investigations. They suffered unnecessary delays in the delivery of much-needed benefits.

Mitigating the risk

Carefully select inputs that are relevant, legal and non-discriminatory. Before making decisions based on correlations, they should be proven to have a cause-and-effect relationship.



02

Transparency and explainability

Transparency is fundamental to both accountability and product optimization. Explainability and interpretability aid in comprehending how decisions are made and provide assurance when AI works well — or recourse when it does not.

Privacy risk

Regulatory enforcement action; intellectual property risks; missing safety features; model design flaws.

What could go wrong

The inability to explain and justify decisions by AI systems protected by 'black box' algorithms prevents individuals from challenging the process and the outcome. Lack of transparency for both the input and output makes it difficult to avoid discriminatory or harmful results.

Mitigating the risk

Personal information to train data sets should be transparent, reliable, accurate, complete and correlated to the data output — which can be challenged if biased or inaccurate. Any data collected should contain privacy notices (right to be informed).

03

Governance and accountability

Privacy laws require robust governance structures and privacy programs that clearly define and communicate roles and responsibilities.

Privacy risk

Concept drift; function creep; flawed algorithm design.

AI developers or project champions can often overstate or even misrepresent claims about their models. When purchasers unquestioningly accept these claims and rely on or greenlight such systems, the shortcomings may only become apparent when real-world harms or other AI failures occur. For example, facial recognition 'false positives' can result in false arrests, detention of innocent people and unwarranted intrusion into their lives.

What could go wrong

A hypothetical self-driving boat designed to find the fastest route across the harbor could potentially cause harm to delicate water systems and unsuspecting swimmers if not programmed to do so safely. So, inputs to training datasets should include images of swimmers.

Mitigating the risk

AI systems require clear lines of accountability to ensure: AI risks are managed effectively; a clear purpose, strategy and set of expectations are communicated to all involved; proper oversight and reporting, especially on deficiencies; third parties (e.g. data providers or AI developers) are held accountable and will cooperate in addressing any issues that arise.



04

Data-minimization

Personal data should be adequate, relevant and limited to the purpose.

Privacy risk

Post-deployment failures (robustness; adversarial attacks; unanticipated interactions).

In the AI world, the common assumption is that 'more is better,' but not all data is of equal quality. Systems trained on data sets without external validation may not function properly in the real world. This can lead to the misuse and unfair re-purposing of someone's personal information. Also, speakers of under-represented languages are likely not consulted before their languages are used to train and develop natural language models without cultural understanding.

What could go wrong

Data sets not externally validated, could result in inappropriate recommendations or conversations with minors, the spread of misinformation and racist remarks, logically inconsistent responses or outright lies. Additionally, using low-quality photos to train facial recognition systems can negatively impact accuracy in real-world applications and create selection bias.

Mitigating the risk

Not all data is useful, relevant or reliable. Data minimization drives better data curation for training, with inappropriate data filtered out, resulting in fewer but higher-quality data sets. Using primarily English and other dominant languages to train LLMs can lead to skewed models and exclude other languages in the training data, especially if only a few media sources (web, social media) are used.

05

Purpose-limitation

Personal data processing should have a clearly defined and communicated purpose to protect rights, respect autonomy and prevent any potential harm.

Privacy risk

Model design and implementation failures (concept drift and data drift through alignment failures); impossible tasks; regulatory enforcement action; outlawed business model.

Lack of purpose opens up the potential for misinformation or disinformation. For example, taking an author's written work and using it to create new articles falsely attributed to the author.

What could go wrong

An LLM trained on publicly available news articles and other data scraped from the web can sometimes provide incorrect answers or promote misinformation or disinformation. If controversial, factually unsound, or defamatory articles are written in the style of an author whose content was used to train the AI model, it could potentially cause harm if the public believes the author wrote it.

Mitigating the risk

All AI system stakeholders, from developers to salespeople and end-users, need to understand and respect the AI's intended purpose. This will guide the selection of data elements used to train the model, the use cases during deployment and operation, the values and assumptions embedded within it, its configuration, safeguards and more. By doing so, the original consent or consent exception and people's expectations are respected, increasing trust and reducing the risk of enforcement activity or public backlash.



06

Accuracy

Under privacy law, personal data must be up-to-date, complete and accurate before it's used. Also, individuals have a right to correct their data.

Privacy risk

Robustness issues (e.g. through overfitting, underfitting); failure under adversarial attacks; unanticipated interactions; missing safety features; unfairness; model drift.

Data quality can impact AI effectiveness, leading to a range of harms. For example, inaccurate data can be detrimental when making decisions about government policy or community planning. Even when data inputs are correct, the model may return an incorrect profile of an individual due to flawed assumptions, poor scoring or inability to process unfamiliar inputs. Inadequate or irrelevant personal data can result in model drift, degrading the model's performance (e.g. lower prediction accuracy). If meaningful human review is lacking, the AI system may make inaccurate decisions, even if it operates as programmed and fails to detect anomalies.

What could go wrong

Police falsely arrest a man because their facial recognition system identifies him as the suspect without other evidence. Despite the police being uncertain about the resemblance between the man and the photo, they believe the accuracy claims of the AI system provider. The man suffers financial harm due to lost pay and legal fees, humiliation, anxiety, inconvenience, loss of liberty and possible stigma due to the wrongful arrest

Mitigating the risk

Data minimization and accuracy principles can help improve data quality to avoid many of these harms. For example, the UK Information Commissioners Office (ICO) advises AI developers to consider the trade-off between data minimization and statistical accuracy at the testing stage to ensure the model remains accurate.⁷ It's essential to monitor model performance for decision-making or predictive AI systems to ensure the data remains relevant, up-to-date, adequate and is retrained where necessary.

⁷ Guidance on AI and data protection, UK Information Commissioners Office, March 15, 2023.

07

Storage-limitation

Privacy laws prohibit businesses from retaining personal data once it's no longer necessary. Still, some laws allow data retention if suitably anonymized.

Privacy risk

Regulatory enforcement action.

Once a model is trained, the underlying training data should generally only be kept if re-training is needed. Even then, the risk of model drift necessitates re-assessing data quality to purge any irrelevant or outdated data. Retaining training datasets beyond their legitimate use, even if they are 'anonymized', also presents a significant compliance burden since re-identification risk must be continuously managed.

Regulators are particularly concerned about breach risks and unlawful re-purposing of datasets or unlawful enrichment and re-identification risk associated with data held longer than necessary. Additionally, even if anonymized, leaked training data can be enriched or reverse-engineered to re-identify individuals.

What could go wrong

A hypothetical AI model designed to assess employment suitability is based on an old dataset prioritizing an employee's ability to attend work in person. The model has not been retrained to address the new realities of remote work. Consequently, candidates unable to attend in person are disadvantaged and excluded or given a lower rating in the hiring process.

Mitigating the risk

Understanding all relevant data retention laws fully and regularly reviewing and purging data where appropriate is vital.



08

Security

Businesses that process personal data must ensure its confidentiality, integrity and availability.

Privacy risk

Failure under adversarial attacks; regulatory enforcement activity; loss of credibility.

Poor security practices can lead to training data being breached, which may include sensitive information like financial details, demographic data and postal codes. Such a breach could expose the people in the training data set to identity fraud risk, financial harm, anxiety and inconvenience in their efforts to avoid potential harms.

What could go wrong

Some key security risks are:

- Re-identification through 'black box' and 'white box' attacks.
- Attribute disclosure risk (risk of inferring additional information from anonymized data).
- Data breaches through adversarial attacks, for example, gaming systems that allow an imposter to gain unauthorized access.

Mitigating the risk

A broader analysis of security will require specialized security expertise. However, some privacy aspects of security (such as the above examples) require specific attention as they impact the full suite of requirements.

09

Respect for end user privacy

AI should respect privacy rights, including rights to information, correction, explanation, deletion and automated decision-making.

Privacy risk

Regulatory enforcement action; loss of credibility and PR headaches; missing safety features.

What could go wrong

Suppose an AI system fails to acknowledge privacy rights and has insufficient checks and balances. In this case, the company exposes itself to regulatory and reputational risk.

Mitigating the risk

These rights apply across the AI lifecycle but may vary slightly by stage. Explainability mechanisms should be built into the entire AI lifecycle, as design, input decisions and modeling can all impact the resulting decision. Review and challenge rights can ensure that people continue to exert control over their personal data. The right to correct input data or challenge assumptions on which a model is based is critical to ensuring fair and accurate decision-making.



Regulatory highlights*

Countries are focusing on overseeing AI's speedy adoption and ensuring that the future includes both business benefits and clear public safeguards for privacy and trust.



Australia has published eight voluntary *Artificial intelligence ethics principles* designed to ensure safe, secure and reliable AI use.

Brazil has set out to establish principles, rules and guidelines to regulate AI development and use. Brazil is considered to be at the forefront of AI policy-making in the region, and given its size and relevance in Latin America, its proposed *AI framework* has the potential to become a regional trendsetter.

Canada's proposed Bill C-27 includes the *Artificial Intelligence and Data Act*, which aims to create new rules for responsible AI development and deployment. The federal government also released a *Companion document* to complement the framework proposed in the act. It's the first step towards a new regulatory system designed to guide AI innovation positively and encourage the responsible adoption of AI technologies by Canadians and Canadian businesses.

China has published its *Internet information service algorithm recommendation management provisions* to safeguard national security and citizens' rights and interests. China has also proposed *Measures for the management of generative artificial intelligence services*.

France has launched a *National strategy for AI* with three key objectives: achieving the highest

scientific level in AI by training and attracting the best global talent, boosting investment in AI and ensuring an ethical approach to AI use and privacy protection. Commission nationale de l'informatique et des libertés (CNIL) has published its *AI resources page*, which includes a detailed self-assessment guide and its *AI action plan page* to prepare for the EU AI Act. The plan also addresses the latest AI developments like GenAI.

India's *Task force on artificial intelligence* in 2018 provided policy recommendations on the ethical deployment of AI that the government can utilize for a five-year period.

Japan's *Social principles of human-centric AI* aims to ensure data privacy and security and create an environment where society can benefit from the data that individuals provide to organizations. Also, the Ministry of Economy, Trade and Industry has published *Governance guidelines for the practice of AI principles*. Emerging AI frameworks will strive to ensure Japan meets its goal of being an AI-ready society.

New Zealand has published the *Treaty of Waitangi/Te Tiriti and Māori ethics guidelines for: AI, algorithms, data and IOT*. These guidelines are intended for government agencies and others engaging with Māori data and communities.

Saudi Arabia in 2020 launched its *National strategy for data and AI*, while its *Personal Data Protection*

Law was amended in 2022 to provide comprehensive data protection during the use of all data relating to individual privacy.

Singapore has introduced its voluntary *AI Verify framework*, which encourages companies to be more transparent about what their AI systems can or cannot do, ideally keeping stakeholders better informed to build trust in AI. The Personal Data Protection Commission (PDPC) recently released its *Model AI governance framework*, along with helpful resources, including the *Implementation and self assessment companion guide for organisations*, *Compendium of use cases*, and an explainer regarding Singapore's *AI Verify Foundation Testing Framework* and *Software Toolkit*. Additionally, the Protection Monetary Authority has published its *Principles to promote fairness, ethics, accountability and transparency (FEAT)* in the use of artificial intelligence and data analytics in Singapore's financial sector.

Spain has announced the first national agency in Europe to supervise AI. With the Agencia Española de Supervisión de la Inteligencia Artificial (AESIA), Spain wants to lead the way in AI regulation in Europe. The Spanish Supervisory Authority (AEPD) recently released *guidelines* for auditing data processing activities involving AI.

South Korea's Personal Information Protection Commission has published its *AI personal information protection self-checklist*.

United Arab Emirates (UAE) has created an *Ethical AI toolkit* to help the industry and the public understand how AI systems can be used responsibly amid the lack of specific legislation governing AI use. In October 2017, the UAE government launched its *UAE strategy for artificial intelligence*.

United Kingdom (UK) has declared appropriate AI use a priority, citing its potential to pose a 'high risk' to individual rights and freedoms. It calls public trust paramount for safe AI adoption. The *National AI strategy* sets out an ambitious 10-year plan for the UK to remain a global AI superpower. The UK proposes a pro-innovation approach to AI regulation. The Information Commissioner's Office (ICO) has provided AI and data protection resources, including the *AI and Data Protection Toolkit*, to aid in this effort. The Centre for Data Ethics and Innovation (CDEI) has also published its *AI Barometer*.

United States (US) has not yet proposed federal AI legislation or regulations. However, some states have enacted state-level AI-related laws, while others are in the process of doing so. The National Institute of Standards and Technology (NIST) has developed an *AI risk management framework* and the Federal Trade Commission (FTC) has issued guidance on *Using artificial intelligence and algorithms*.

* These highlights are not an exhaustive list but merely illustrative of the current developments that are in flight at the time of writing.



The road ahead: Building trustworthy AI

Data privacy is the bedrock of an AI-first enterprise, with transparency over how personal information is deployed — and full accountability for any misuse, with processes for fast mitigation.

AI's speed and efficiency are transforming the world, and companies are understandably keen to harness its potential. Technological advantage can only bring competitive rewards if customers and other stakeholders trust that data is being used responsibly.

While regulators are still struggling to keep up with AI advancements, the EU AI Act sends a strong message that regulation will be comprehensive with significant consequences for non-compliance. Businesses that build robust controls over AI use based on clear ethical guidelines should be well-positioned to leverage AI's benefits while safeguarding society from potential risks and satisfying regulatory requirements.

Five key steps that can help companies build trust in AI.

1 Understand your regulatory environment and adopt an AI privacy strategy

Legislators, policymakers and regulators consistently stress aligning AI systems with recognized standards. So, it's essential to identify which regulatory frameworks apply to your business, determine which you choose to comply with and plan how your AI will be deployed. Create a baseline for AI usage that satisfies varying regimes and streamline your AI development or AI-related business activities accordingly.

2

Incorporate Privacy by Design into your AI projects

Assess the impact on privacy and address compliance issues at the ideation stage — and throughout the AI lifecycle — through a systematic privacy impact assessment (PIA) or data protection impact assessment (DPIA). Privacy by Design, as outlined in the ISO 31700 Privacy by Design Standard and KPMG's Privacy by Design Assessment Framework, can help organizations build privacy into AI systems.

Even if you believe your system only uses anonymized or non-personal data, privacy risks can emerge, including re-identification from training data sets and even AI models and downstream impacts of non-personal data used to train models that impact individuals and communities. A robust assessment will also include security and privacy threat modeling across the AI lifecycle and stakeholder consultation where appropriate. Consider broader privacy issues such as data justice (how fairly people are treated in the way you use their data) and indigenous data sovereignty (the rights of indigenous peoples to govern data about their communities, peoples, lands and resources).

3

Assess AI privacy risks

Either in conjunction with or after your PIA/DPIA, assess AI privacy risks. Be sure to adhere to major standards, best practices and codes of conduct to operationalize the requirements. This applies whether you are the developer or a client acquiring and integrating an AI system. If you are a client, ask the developer for documentation to support their PIA and related AI privacy risk assessments. If they can't provide this documentation, consider another provider. In many jurisdictions, including the UK and the EU, a PIA/DPIA is a legal requirement.

4

Audit your AI system

If you are a developer of AI systems or a third party/vendor of AI, you should assure clients and regulators that you have taken the necessary care to build trustworthy AI. One way to do this is through an audit against recognized standards, regulatory frameworks and best practices, including an algorithmic impact assessment.

5

Respect rights and choices through explainability and transparency about data inputs and outputs

Be prepared to answer questions and manage the preferences of individuals impacted by your development or use of AI systems. Organizations that want to use AI for automated decision-making should be able to explain in plain language how AI will impact their end users. Explainability is the capacity to articulate why an AI system reached a particular decision, recommendation or prediction. Be prepared to answer questions and manage the preferences of individuals impacted by your development or use of AI systems. Consider developing documented workflows to identify and explain what data was used, how it was applied to the end user and how the end user can contest or challenge the use of AI for decision-making purposes.



How this connects with what we do

With roots stretching back over 150 years, KPMG firms have played a leading role in exploring and harnessing new technologies, such as GenAI, and providing assurance and direction in implementing them.

We understand that responsible AI is a complex business, regulatory and technical challenge. KPMG firms are committed to helping clients bring a responsible AI offering to life. Using GenAI responsibly, KPMG helps organizations build trustworthy and safe AI tech solutions.

Further, KPMG privacy risk management professionals take a responsible approach to assessing the ethics, governance and security in place around

clients' AI and machine learning technologies. We aim to give you a holistic view of privacy trends and offer a range of privacy solutions and services, with expertise in PrivaTech, Privacy by Design, privacy target operating models, ESG and privacy, privacy program design, implementation and automation, customer experience and change management, privacy breach response and privacy remediation.

KPMG firms are helping businesses in every sector embrace a new era of opportunity in the digital economy. From strategy to implementation, KPMG professionals can help transform your current business model to drive future competitiveness, growth and value. KPMG. Make the Difference



KPMG Connected Enterprise

KPMG's customer centric, agile approach to digital transformation, tailored by sector.



KPMG Powered Enterprise

KPMG's suite of services to transform functions. Target operating models designed with the future in mind, using KPMG leading practices and processes and pre-configured SaaS (Software as a Service) platforms.



KPMG Trusted

How to build and sustain the trust of your stakeholders.



KPMG Elevate

Unlock financial value quickly and confidently.



Our lead author



Sylvia Klasovec Kingsmill
Global Cyber Privacy Leader
KPMG International and Partner
KPMG in Canada

Sylvia is the Global Cyber Privacy Leader with over 20 years of experience in risk management for both public and private sectors. She is based in Toronto, Canada, and helps executive teams implement data-driven and AI strategies into their business transformations while ensuring adherence to privacy regulations, policies and AI governance trends. Sylvia specializes in privacy tech enablement and automation, helping clients integrate holistic privacy solutions throughout their operating environments and functions.

In 2022, Sylvia was recognized as one of the top 10 Canadian privacy leaders by a leading law firm. She also launched an ISO 31700 Privacy by Design Certification Program to assess emerging tech and demonstrate compliance with international data protection laws and privacy best practices. Sylvia has advised on some of the biggest data breaches and has extensive experience transforming corporate compliance and privacy programs in response to internal audits, regulatory findings and compliance orders.

Acknowledgements

This report would not be possible without the invaluable planning, analysis, writing and production contributions of colleagues around the world.

Global contributors

Abigail Dubiniecki, Freelance Privacy Consultant

Andrew Sprague, KPMG in Canada

David Ferbrache, KPMG International

Billy Lawrence, KPMG International

Leonidas Lykos, KPMG International

Global collaborators

Paul Henninger

Global Lighthouse Leader
KPMG International and Partner
KPMG in the UK
paul.henninger@kpmg.co.uk

Leanne Allen

Data, Data Science and AI Lead
Partner, KPMG in the UK
leanne.allen@kpmg.co.uk

Isabel Simpson

Head of Data Privacy
Partner, KPMG Law in the UK
Isabel.simpson@kpmg.co.uk

Kelly Henney

Privacy and Data Protection Services
Leader
Partner, KPMG Australia
khenney@kpmg.com.au

Orson Lucas

Data Protection and Privacy Services
Leader
Principal, KPMG in the US
olucas@kpmg.com

Christophe Buschmann

Data Protection Lead
Director, KPMG in Luxembourg
christophe.buschmann@kpmg.lu



Our contacts

Sylvia Klasovec Kingsmill

Global Cyber Privacy Leader
KPMG International and Partner
KPMG in Canada
skingsmill@kpmg.ca

Walter Risi

KPMG in Argentina
wrisi@kpmg.com.ar

Kelly Henney

KPMG Australia
khenney@kpmg.com.au

Veronica Scott

KPMG Law in Australia
vscott1@kpmg.com.au

Andreas Tomek

KPMG in Austria
atomek@kpmg.at

Dieter Buchberger

KPMG Law in Austria
dbuchberger@kpmg-law.at

Jeyapriya Partiban

KPMG in Bahrain
jeyapriyapartiban@kpmg.com

Allison James

KPMG in Barbados and Eastern
Caribbean
allisonjames@kpmg.bb

Benny Bogaerts

KPMG in Belgium
bbogaerts@kpmg.com

Leandro Augusto M Antonio

KPMG in Brazil
lantonio@kpmg.com.br

Erick Palencia

KPMG in Chile
erickpalencia@kpmg.com

Andres Martinez

KPMG Law in Chile
avmartinez@kpmg.com

Henry Shek

KPMG China
henry.shek@kpmg.com

Felipe Silgado

KPMG in Colombia
fsilgado@kpmg.com

Bryan Beesley

KPMG in the Crown
Dependencies (Guernsey,
Isle of Man, Jersey)
bbeesley@kpmg.co.im

Pavel Kliment

KPMG in the Czech Republic
pkliment@kpmg.cz

Martin Hrdlik

KPMG Law in the Czech Republic
mhrdlik@kpmg.cz

Urmass Kukk

KPMG Law in Estonia
urmas.kukk@kpmglaw.ee

Tomi Nevasoja

KPMG in Finland
tomi.nevasoja@kpmg.fi

Ari Engblom

KPMG Law in Finland
ari.engblom@kpmg.fi

Vincent Maret

KPMG in France
vmaret@kpmg.fr

Julie Bellesort

KPMG Law in France
jbellesort@kpmgavocats.fr

Jaba Gvelebiani

KPMG Law in Georgia
jgvelebiani@kpmg.com

Michael Falk

KPMG in Germany
mfalk@kpmg.com

Francois Maartens Heynike

KPMG Law in Germany
fheynike@kpmg-law.com

Theodossis Tompras

CPA Law (Greece)
ttompras@cpalaw.gr

Bálint Tóásó

KPMG Law in Hungary
balint.toaso@kpmg.hu

Mayuran Palanisamy

KPMG in India
mpalanisamy@kpmg.com

Freddie Mulyadi

KPMG in Indonesia
freddie.mulyadi@kpmg.co.id

Tom Hyland

KPMG in Ireland
tom.hyland@kpmg.ie

Lynn Brill-Sarusi

KPMG in Israel
lbrillsarusi@kpmg.com

Luca Boselli

KPMG in Italy
lboselli@kpmg.it

Ravi Sankar

KPMG in Jamaica
rsankar@kpmg.com.jm

Manabu Katsumura

KPMG in Japan
manabu.katsumura@jp.kpmg.com

Nancy Mosa

KPMG in Kenya
nmosa@kpmg.co.ke

Min Soo Kim

KPMG in Korea
mkim9@kr.kpmg.com

Una Petrauska

KPMG Law in Latvia
upetrauska@kpmglaw.lv

Christophe Buschmann

KPMG in Luxembourg
christophe.buschmann@kpmg.lu

Albert Lim

KPMG in Malaysia
hockenglim@kpmg.com.my

Rommel Garcia

KPMG in Mexico
rommelgarcia@kpmg.com.mx

Koos Wolters

KPMG in the Netherlands
koos.wolters@kpmg.nl

Philip Whitmore

KPMG in New Zealand
pwhitmore@kpmg.co.nz

John Anyanwu

KPMG in Nigeria
john.anyanwu@ng.kpmg.com

Tor Henning Rustan Knudsen

KPMG Law in Norway
tor.knudsen@kpmg.no

Jallain Manrique

KPMG in the Philippines
jsmanrique@kpmg.com

Michał Kurek

KPMG in Poland
michalkurek@kpmg.pl

Magdalena Bęza

KPMG Law in Poland
mbeza@kpmg.pl

Sérgio Martins

KPMG in Portugal
smmartins@kpmg.com

Gabriel Mihai Tanase

KPMG in Romania
mtanase@kpmg.com

George Manea

KPMG Law in Romania
georgemanea@kpmg.com

Ton Diemont

KPMG in Saudi Arabia
antondiemont@kpmg.com

Marija Milojevic

Milojevic & Zlatanovic,
Legal partnership (Serbia)
mmilojevic@kpmg.com

Wendy Lim

KPMG in Singapore
wlim@kpmg.com.sg

Pavol Adamec

KPMG in Slovakia
padamec@kpmg.sk

Matjaž Pušnik

KPMG in Slovenia
matjaz.pusnik@kpmg.si

Finn Elliot

KPMG in South Africa
finn.elliott@kpmg.co.za

Javier Aznar Garcia

KPMG in Spain
jaznar@kpmg.es

Daniel Grantz

KPMG in Sweden
daniel.grantz@kpmg.se

Mattias Bergman

KPMG Law in Sweden
mattias.bergman@kpmg.se

Thomas Bolliger

KPMG in Switzerland
tbolliger@kpmg.com

Jason Y.T. Hsieh

KPMG in Taiwan
jasonhsieh@kpmg.com.tw

Sagulyat Naipaporn

KPMG in Thailand
naipaporn@kpmg.co.th

Onur Küçük

KP Law (Turkey)
onur.kucuk@kplawtr.com

Timothy Wood

KPMG in the UAE
timothywood@kpmg.com

Alex Yankovski

KPMG in Ukraine
ayankovski@kpmg.ua

Graham Thomas

KPMG in the UK
graham.thomas@kpmg.co.uk

Isabel Simpson

KPMG Law in the UK
Isabel.simpson@kpmg.co.uk

Rodrigo Ribeiro

KPMG in Uruguay
rribeiro@kpmg.com

Orson Lucas

KPMG in the US
olucas@kpmg.com

Hanh Linh Truong

KPMG in Vietnam
linhtruong@kpmg.com.vn

Amarjit Kaur

KPMG Law in Vietnam
amarjitsingh@kpmg.com.vn



Some or all of the services described herein may not be permissible for KPMG audit clients and their affiliates or related entities.

kpmg.com



The information contained herein is of a general nature and is not intended to address the circumstances of any particular individual or entity. Although we endeavor to provide accurate and timely information, there can be no guarantee that such information is accurate as of the date it is received or that it will continue to be accurate in the future. No one should act on such information without appropriate professional advice after a thorough examination of the particular situation.

© 2023 Copyright owned by one or more of the KPMG International entities. KPMG International entities provide no services to clients. All rights reserved.

KPMG refers to the global organization or to one or more of the member firms of KPMG International Limited ("KPMG International"), each of which is a separate legal entity. KPMG International Limited is a private English company limited by guarantee and does not provide services to clients. For more detail about our structure please visit kpmg.com/governance.

The KPMG name and logo are trademarks used under license by the independent member firms of the KPMG global organization.

Throughout this document, "we," "KPMG," "us" and "our" refers to the global organization or to one or more of the member firms of KPMG International Limited ("KPMG International"), each of which is a separate legal entity.

Designed by Evalueserve.

Publication name: Privacy in the new world of AI | Publication number: 138971-G | Publication date: September 2023