Deloitte.

How to unlock the power of AI and protect privacy

A must-have strategy for every organisation

Artificial Intelligence (AI) is making rapid strides in the corporate world, unlocking potential opportunities for efficiency, automation, and problem-solving. But alongside the benefits, come privacy and data security concerns. Why? AI systems use vast amounts of data to learn and make predictions, so the potential for bias, security issues, loss of control and privacy breaches has become a real and pressing issue for every organisation.

There are, however, innovative ways organisations can mitigate challenges to unleash the power of AI.

Introduction

From personalised recommendations to medical diagnoses and smart cities, AI technologies are transforming our lives and reshaping industries. But with AI's meteoric rise, come concerns about data privacy and security. With AI systems heavily relying on vast troves of data for training and analysis, individuals' personal and sensitive information becomes susceptible to potential breaches and misuse.

In response to these challenges, the field of privacy preserving AI (PPAI) has emerged, aiming to strike a delicate balance between harnessing the power of AI and safeguarding individuals' privacy rights. PPAI encompasses a diverse range of methodologies, protocols, and algorithms that enable AI applications to glean insights from data without directly retrieving or compromising personal and sensitive information.

Artificial intelligence defined

AI refers to the simulation of human intelligence in machines programmed to perform tasks, such as learning, reasoning, problem-solving, and decision-making. AI encompasses a wide range of technologies, including machine learning, natural language processing, computer vision, and robotics. The goal of AI is to enable machines to process and analyse vast amounts of data, recognise patterns, and adapt their behaviour accordingly, leading to more efficient and intelligent systems.

Benefits of AI

AI's benefits for modern life are diverse and far-reaching. AI enhances efficiency and productivity by automating repetitive tasks and streamlining complex processes across industries. It enables personalised experiences, from tailored recommendations to targeted marketing, improving customer satisfaction and engagement. AI-driven insights aid in data analysis, decision-making, and risk assessment, empowering businesses to make informed strategic choices. In healthcare, AI applications facilitate early disease detection, drug discovery, and precision medicine, revolutionising patient care.

With its continuous evolution, AI promises to unlock even greater opportunities, driving innovation and enriching the human experience.

Key concerns around personal information and AI

Using personal information to develop and test AI has raised significant ethical and privacy concerns. AI systems require vast amounts of data, including, in some circumstances, the use of personal information, for training and fine-tuning. While using this information can lead to powerful and accurate AI models, it also poses risks of data breaches, identity theft, and misuse of sensitive information. There are several issues associated with the interaction between AI and personal information. Some key concerns are:

Privacy breaches

AI systems may process or store personal information and if not handled properly, there's a risk of unauthorised access or loss of data. If personal information falls into the wrong hands, it can lead to identity theft, fraud or other privacy violations.

Profiling and manipulation

AI systems that extensively profile individuals based on personal information can potentially lead to manipulation or discriminatory practices.

Re-identification and de-anonymisation

Even if personal information is anonymised, there's a risk of re-identification or de-anonymisation. By combining seemingly anonymous data with other information sources, it may be possible to link it back to specific individuals, compromising their privacy.

Lack of control

Individuals often have limited control over their personal information once it's processed by AI systems. The data may be stored indefinitely, shared with third parties or used for purposes beyond the original scope. This lack of control can undermine individuals' privacy and autonomy.

Ensuring responsible data-handling practices and anonymising data are essential steps towards preserving privacy rights during AI development. Striking a balance between leveraging personal information for AI advancements and protecting individuals' privacy remains a critical challenge that necessitates robust data protection regulations and ethical guidelines in the ever-expanding AI landscape.

An innovative approach: privacy preserving AI (PPAI)

PPAI aims to reconcile AI's immense potential while safeguarding individual privacy. It involves a range of techniques and algorithms designed to extract meaningful insights from data while ensuring anonymity and confidentiality of personal and sensitive information.

Examples of PPAI capabilities

Fully homomorphic encryption (FHE)

FHE is a special type of encryption allowing complex calculations to be performed on encrypted data without needing to first decrypt it. This enables AI models to be trained on encrypted data, eliminating the risk associated with decryption, thereby ensuring enhanced data privacy and security.

Example: a bank might use AI to analyse customer spending habits to offer personalised financial advice. However, customer financial data is highly sensitive. With FHE, the bank can encrypt this data and then use AI to analyse it while still encrypted. The AI can identify spending patterns and give advice without ever seeing the actual data. This means the bank can offer valuable, customised services to its customers while ensuring their financial information remains completely secure and private.

Federated learning

Federated learning is a distributed machine learning approach where the learning process is decentralised across multiple devices (rather than one central place, like a server), this allows the model to learn from data without centralising or storing it.

Example: a healthcare application (app) sends a learning model to each hospital. This model learns from the hospital's data, but the data never leaves the hospital. Instead, only the learned information (like patterns or insights) is sent back to the app. The app combines these insights from all hospitals to improve its ability to diagnose diseases or suggest treatments. This way, the app gets smarter without compromising patient privacy.

An innovative approach: privacy preserving AI (PPAI)

Secure multiparty computation (MPC)
MPC is a technique allowing multiple parties to jointly compute a function over their individual private inputs without revealing inputs to each other. The parties can collaboratively perform computations while keeping their individual data confidential.

Example: three hospitals want to jointly analyse patient data to develop a medical model without sharing individual patient records. Using MPC, they can compute statistical measures, such as averages and correlations, on their respective datasets without exposing the raw data to each other. This ensures data privacy while enabling collaborative analysis.

Trusted execution environment (TEE)
TEE is a secure and isolated environment within a computer system in which sensitive computations can be performed with strong guarantees of privacy and security. TEEs ensure code and data are protected from tampering, even by the operating system or other software.

Example: financial data is encrypted and sent to the cloud. Once in the cloud, it's moved into the TEE, a highly secure area where no one, not even the cloud provider, can peek inside. The AI algorithms run inside this enclave, analysing the data securely. After processing, the results are sent back to the company, still encrypted. This way, the company can leverage powerful cloud-based AI without compromising the confidentiality of its data.

Synthetic data
Synthetic data is like a realistic copy of real-world data, created using algorithms and computer simulations. It's designed to mimic the patterns and characteristics of actual data without containing any real, sensitive information.

Example: a retail company wants to improve its recommendation system using AI but doesn't have enough customer shopping data. Instead of waiting to collect more real data, the company uses synthetic data. They create a large, artificial dataset that resembles real customer shopping behaviours but doesn't contain any actual customer information. The AI is then trained on this synthetic data. It learns patterns and preferences just like it would from real data, allowing the company to enhance its recommendation system.

An innovative approach: privacy preserving AI (PPAI)

Differential privacy
Differential privacy works by adding noise to the data in a controlled way. This noise is small enough not to significantly affect the accuracy of the analysis, but large enough to prevent individual records being identified.

Example: advertisers want to show users relevant ads while also protecting users' privacy. Differential privacy can train a machine learning model to predict which ads are most relevant to each user without revealing individuals' data.

Applications and benefits

When it comes to AI development and deployment, PPAI offers many benefits and addresses many critical data privacy and security concerns.

Protects individual privacy
AI can be used to collect and analyse data from a variety of sources, including medical records, financial data, and social media data. PPAI lets organisations collect and analyse data without revealing individual users' sensitive information, therefore protecting individuals' privacy while still benefitting from AI insights.

Enables new forms of collaboration
PPAI enables new forms of collaboration between organisations without requiring them to share their data with each other, allowing them to pool resources and expertise to develop new AI models and solve complex problems.

Builds trust with users
When users know their privacy is being protected, they're more likely to trust organisations and use their products and services. PPAI can help organisations build trust with users by demonstrating they're committed to protecting their privacy.

Protect intellectual property rights
Innovative AI models can be encrypted and deployed in third party environments without exposing the AI's algorithms. This ensures any novel algorithms and coding aren't copied without permission.

Overall, PPAI empowers responsible and ethical use of data, paving the way for a privacy-conscious and secure AI ecosystem. It fosters collaborations and data sharing without the need to disclose raw sensitive information, encouraging collective intelligence and knowledge exchange.

Empowering innovation

We help organisations assess their PPAI maturity, identify privacy risks, and develop PPAI roadmaps. Additionally, we can support in the design, development, and integration of privacy-preserving techniques into existing AI systems.

By partnering with Deloitte, you can be confident you're taking the right steps to protect the privacy of your customers and employee's information while still benefiting from the power of AI.

Conclusion

PPAI is a rapidly developing field with the potential to revolutionise how we use AI.

By developing and using PPAI techniques, we can help protect people's data, reduce AI risks, and enable new forms of innovation.

As AI becomes more integrated into our lives, it's vital safeguards are in place to protect our privacy. PPAI helps us use AI in a safe and responsible way by collecting and analysing data without revealing individuals' sensitive information.

PPAI isn't just a technological innovation, it represents a moral imperative. Deloitte is committed to guiding organisations towards a future where innovation, privacy protection and cyber security go hand in hand.

Contacts

Daniella Kafouris
Partner
dakafouris@deloitte.com.au

Kate Monckton
Partner
kmonckton@deloitte.com.au

Tim Scott
Director
timscott@deloitte.com.au

Galen Ou
Senior Analyst
gou@deloitte.com.au

Deloitte.

1256910258_Designed and produced by The Agency | Deloitte Australia_12/23