Data privacy in the age of Artificial Intelligence: Challenges and solutions

Information technology is undergoing the Big Bang, in which the size of data generated worldwide doubles every two years. This unparalleled age of digital breakthroughs is propelled by billions of smartphones and other devices. The volume, variety, and velocity of data generation have led to a paradigm shift in data-driven decision-making.

Transcending human intelligence, increasingly powerful and sophisticated software tools are taking over decision-making in business environments. Businesses harness the power of cutting-edge artificial intelligence (AI) technologies to draw insightful inferences and make predictions about user behavior.

AI technologies mainly rely on data fed into the system to train models and improve their performance. Most often, this data comprises sensitive personal data about users, with which businesses aim to personalize user experiences and optimize targeted advertising purposes. Extensive use of AI, however, raises crucial concerns about user privacy.

In this blog post, we will try to decode data privacy in the age of AI by understanding its challenges and solutions.

Privacy challenges in AI and their respective solutions

The problem of Re Identification

AI is mainly known for seemingly infinite pattern recognition capabilities and establishing connections between various data points, which pose a risk of unauthorized reidentification. Many studies have highlighted how the information that is anonymized and scrubbed of all identifiers can still be reidentified using emerging computational strategies.

When different data sets (with no trace of personally identifiable information present in any of them) are combined, it gives birth to the ‘Data Mosaic effect’, which is the ability to uniquely identify an individual. It increases the privacy risks of allowing private AI companies to process the personal data of consumers, even in circumstances where anonymization is carried out.

Pseudonymization techniques like masking of direct identifiers were once considered sufficient to “anonymize” a dataset. However, recent developments in linkage attacks, which are performed by employing data mining and machine learning to look for overlapping matches between the common attributes, or quasi-identifiers, of two or more data sets, have exposed the vulnerabilities with pseudonymization to protect privacy.

Profiling attacks constitute another kind of privacy attack that leverages AI to re-identify individuals based on their behavioral patterns. In peer-reviewed papers, for example, joint research by the Vienna University of Economics and Business and MOSTLY AI in their paper, demonstrates how profiling attacks are possible to successfully re-identify browsing patterns of individuals.

Solution

The unauthorized risk of reidentification by AI can be effectively mitigated by embracing technologies that allow organizations to realize commercial objectives without compromising data privacy and security. Advancements in privacy-preserving techniques like differential privacy and synthetic data generation offer promising alternatives.

Organizations that prioritize the value of consumer trust avoid transferring actual production data (the final stage of a product’s lifecycle, readily accessible to end users, whether through websites or mobile apps) into non-production (development and testing) environments. They understand the principle of data minimization and that the data of a customer should be utilized to serve that actual customer only.

Hence, they adopt statistically representative synthetic data for privacy-preserving AI and analytics. Synthetic data generation (SDG) involves creating artificial data sets with similar statistical properties but without any trace of identifiable information about real individuals. SDG eliminates the one-to-one correlation between data points and actual customers, making it significantly cumbersome to re-identify individuals.

The Power of Consent

Most existing privacy laws, including Federal Trade Commission enforcement, rely on a model of consumer choice built around ‘notice-and-consent’. In an AI-driven world, the ‘notice-and-choice’ model may not be deemed fit. The current consent models rarely offer fine-grained control mechanisms for AI purposes.

As per GDPR Article 6(1)(a), even if users consent to their data being collected, it is doubtful that the consent would be truly "informed." They likely may not be able to comprehend the extent of unpredictable AI uses, given that the technology rapidly continues to evolve and may involve unforeseen data uses over time.

Furthermore, the requirements for the contract as a legal basis for AI processing of EU personal data under GDPR Article 6(1)(b) will not be satisfied; there are more less-intrusive, "data minimized" means of carrying out the performance of the contract without requiring AI processing.

Consent models also do not fit well for applications with critical infrastructure like self-driving cars and smart traffic signals. For example, imagine a self-driving car pausing mid-intersection for a driver to review and accept data permission used for split-second decision-making, like a sudden maneuver to avoid an accident. Such real-time decisions require immediate action and can't wait for one's approval, rendering the notice-and-consent model ineffective.

Solution

Data Value Proposition and Transparency

With a data value proposition, a company showcases what value it returns in exchange for user data. It could be personalized recommendations, enhancing the user experience, or having users have a say in how their data is used. Users (as much as 57%) likely feel okay trading their data in exchange for discounts or offers.

In a growing privacy-awareness environment where concern about data misutilization makes users avoid websites that ask for personal information, having a strong data value proposition that addresses these concerns and convinces users that their data is used responsibly makes more sense than ever for data-driven companies.

Degree of Personalization

Users should have the freedom to choose the level of personalization they desire to be generated by AI systems. Given that users have the choice to choose how their data is used and for what purposes, the solution may address the personalization-privacy paradox. Whilst notice and consent model has been argued to solve this paradox, the binary nature of this approach does not solve the preferences and comfort unique to each customer.

Therefore, personalization as a continuum would provide a clearer picture. It can accommodate individual preferences by offering granular control over the level of personalization desired. For example, customers may choose to opt for:

- No personalization
- Minimal personalization (e.g., newsletters)
- Moderate personalization (personalized content and segmentation)
- Complete personalization (e.g., third-party integration)

Alternatives for consent and contract

If both consent and contract prove insufficient as lawful grounds for AI processing, legitimate interest under GDPR Article 6(1)(f) will most likely act as an alternative for processing in compliance with the EU AI Act, the proposed legal framework governing the use of AI in the EU. When necessary for a compelling public interest, like safety in self-driving cars, legitimate interest empowers organizations to process data without obtaining the explicit consent of users.

Issues of biases and discriminations

Algorithmic bias, or AI bias, refers to the tendency of algorithms to reflect human biases. Machine learning relies on data fed into the system to train the models; this data often comprises the personal data of consumers.

If the data fed into the system contains biases, the resulting algorithms will also tend to be biased. Because these systems are trained to replicate the outcomes derived by human decision-makers, AI may potentially perpetuate existing inequalities and discrimination.

If biases are manifested in the data to represent one group over another, the algorithm may learn to associate features with specific groups. For example, if a facial recognition algorithm is trained to easily identify a white person compared to a black person, it will unintentionally hinder equal opportunity for people from those specific groups.

Solution

Addressing bias in AI requires combining technical solutions with social and policy measures. Inclusion of responsible practices and prioritizing fairness throughout the AI lifecycle can aid in building a bias-free AI world without amplifying existing disparities.

Data-driven approach

- Diverse teams comprising data scientists and engineers, social scientists and ethicists, legal and policy experts, as well as domain experts, should proactively identify and remove biased data points before training the models with them.
- Use of techniques like data generation or oversampling for underrepresented groups can increase their representation in the training data and balance out skewed demographics.

Algorithmic and technical approach

- Transparency in algorithms should be maintained by the implementation of explainable AI models. Explainable AI models allow users to comprehend how AI decisions are made and identify potential bias in the reasoning process.

- Techniques like counterfactual fairness assessments help determine how different individuals or groups would be treated under the AI system, potentially uncovering discriminatory possibilities.
- AI systems should be continuously monitored for bias and discrimination, in addition to human intervention and corrective measures where required.

Legal approach

In the US, privacy stakeholders have expressed desires to change the current 'consent-and-choice' model to shift the privacy protection responsibility from consumers over to the businesses that collect and process personal data. The proposed model focuses on regulating companies' processing of data—what they collect and how they can use or share it.

In AI-driven decision-making, this model can address issues that arise due to any algorithmic discrimination in several ways, including:

- Data stewardship requirements, such as robust enforcement and accountability and duties of fairness or loyalty, could oppose the adverse uses of individuals' personal data.
- Data transparency or disclosure measures like algorithmic auditing and explainability tools could identify and address biases in AI-driven decision-making.
- Implementation of data governance structures, including privacy officers, impact assessments, and privacy-by-design principles, may highlight issues related to the use of algorithms.
- Rules on data collection and sharing could reduce data aggregation that enables discriminatory inferences.

Final Words

AI, without a doubt, is the future. Stepping into a future that respects the privacy of users while benefiting businesses simultaneously could be done by focusing on user-centric data practices, promoting responsible AI developments, and fostering a public-private collaborative ecosystem.

Algorithms' potential to be biased when weighing people or re-identifying data requires thorough understanding and addressing the partiality and vulnerabilities first before developing solutions to create unprejudiced AI systems.

As AI continues to evolve, it is imperative for businesses and regulators alike to remain vigilant in addressing these challenges, ensuring the technology is used for greater good rather than having an adverse impact on consumers' right to privacy. In an ever-evolving

digital landscape, a safe way forward is to balance the potential benefits of AI with the crucial need to protect user privacy.