The impact of privacy-preserving technology on data protection

In an era dominated by digital advancements, safeguarding personal data has become paramount. The escalating frequency of data breaches and privacy concerns has prompted a revolutionary response from the tech community – Privacy-Preserving Technologies (PPTs). This article delves into the transformative impact these technologies have on data protection, exploring their significance, methodologies, and the resulting changes in the privacy landscape.

In this article:

The Current Landscape of Data Protection

Before exploring the impact of PPTs, it's crucial to acknowledge the existing challenges in data protection. Data breaches, privacy concerns, and the dynamic landscape of regulatory compliance underscore the urgent need for advanced solutions.

Examining the medical sector provides a striking example of these challenges. In an unsettling experiment conducted by the Mayo Clinic, imaging and facial recognition technologies were employed to match research subjects with their anonymized brain scans.[1] Reconstruction of facial features from brain scans allowed a facial recognition program to accurately match 70 out of 84 subjects.[2] This highlights the vulnerability of even anonymized medical data to potential deanonymization.

Anonymization attempts, such as "de-facing" brain scans, are not foolproof, as neural networks can reconstruct facial features and compromise anonymity. Even for datasets without faces or images, anonymization may not be a robust solution.[3] True anonymization is elusive, as data can be combined with other public datasets, leading to potential deanonymization (see the Netflix deanonymization in our article “The Most Common Data Anonymization Techniques”). Security expert Bruce Schneier notes that seemingly anonymous data, such as gender, date of birth, and city, can still make about half of the US population identifiable.[4]

---

[1] Privacy Preserving Tech, Tools for data use.
https://blog.openmined.org/privacy-preserving-tech-tools-for-safe-data-use/
[2] The New York Times, You Got a Brain Scan at the Hospital. Someday a Computer May Use It to Identify You.
https://www.nytimes.com/2019/10/23/health/brain-scans-personal-identity.html
[3] https://arxiv.org/pdf/1810.06455.pdf
[4] https://www.schneier.com/essays/archives/2007/12/why_anonymous_data_s.html

Moreover, researchers at Oxford have demonstrated that age and sex can be inferred from structural brain images or directly identified by heartbeat patterns.[5] Once health data is exposed, it becomes public, carrying significant consequences.

Data encryption is considered as an alternative to anonymization. However, this introduces its own set of challenges. While encryption safeguards sensitive data, decrypting it is often necessary for training sophisticated AI models. Health devices, smart energy devices, and assisted driving systems, for instance, collect vast amounts of sensitive data that may need to be decrypted for model training. This poses risks of private data exposure to those working on the models, potential inference attacks on the resulting model, and, finally, potential loss or theft of unencrypted data.

In the era of AI, ensuring data safety is paramount. Today's mathematical capabilities enable the reverse engineering of even anonymized datasets, leading to the identification of individuals and inferences about their private lives. In response to these challenges, PPTs emerge as a paradigm shift in how we handle and protect sensitive information.

The Rise of PPTs

Understanding PPTs

PPTs are a category of tools, methodologies, and techniques designed to protect the privacy of individuals and sensitive information in various digital interactions. These technologies aim to enable the secure handling, sharing, and processing of data while minimizing the risk of unauthorized access, disclosure, or compromise. PPTs are particularly important in the context of increasing concerns about data breaches, cyber threats, and the need to comply with privacy regulations.

PPTs encompass a range of innovative approaches designed to secure sensitive information without compromising its utility. Among these, differential privacy, federated learning, secure multiparty computation, and homomorphic encryption stand out.

We'll take a quick look at three of them before discussing Differential Privacy forward:

- **Federated Learning** is a machine learning approach that allows a model to be trained across multiple decentralized edge devices or servers holding local data samples without exchanging them. The concept behind federated learning is to enable collaborative model training without centrally pooling the raw data, thereby addressing privacy concerns and reducing the need to transfer sensitive information to a central server.

  As we will see forward, federated learning has practical applications in various domains, including healthcare (for patient data privacy), finance (for fraud detection without centralizing transaction details), and edge computing scenarios (where training on local devices is preferred due to latency or bandwidth constraints). While it introduces challenges such as communication overhead and ensuring the quality of

[5] https://www.technologyreview.com/2019/06/27/238884/the-pentagon-has-a-laser-that-can-identify-people-from-a-distanceby-their-heartbeat/

model updates, federated learning continues to be an evolving and promising approach in the field of machine learning.

- **Secure Multiparty Computation (SMPC)** is a cryptographic technique that enables multiple parties to jointly compute a function over their inputs while keeping those inputs private. The fundamental concept behind SMPC is to distribute the computation across multiple parties in such a way that none of them can learn anything beyond the final output. This is achieved using cryptographic protocols that enable secure interactions among the parties. SMPC ensures that participants can collectively compute a result based on their combined inputs without compromising the privacy of the raw data.

  While SMPC offers robust privacy guarantees, it also comes with challenges such as communication overhead, computational complexity, and the need for well-designed protocols. Advances in cryptography continue to refine and expand the applications of secure multiparty computation in various domains.

- **Homomorphic Encryption** is a cryptographic technique that allows computations to be performed on encrypted data without decrypting it. In other words, it enables data to remain encrypted while mathematical operations are carried out on the encrypted data. The result of these operations, when decrypted, matches the result of performing the same operations on the original, unencrypted data. For example, it would make it possible to find data on people with arthritis from a wearables data set, run calculations on it, and create a useful model based on group-level insights without ever decrypting personal records. Homomorphic encryption is gaining popularity, and it is hoped that one day, almost all computation will be done on encrypted data.

What is Differential Privacy and How Does it Work?

Today, differential privacy (DP) is widely recognized as the gold standard in the field of anonymization, allowing for robust statistical analysis without compromising the privacy of the dataset. Tech giants like Google and Apple are using DP for many of their applications, giving it the "stamp of approval." Specifically, DP enables data scientists to query a database while providing guarantees about the privacy of the records within it. To illustrate this concept, consider a database with rows representing individuals and columns indicating binary values (0 or 1) for various attributes, such as gender or the presence of a disease—information worth safeguarding.

Example Database	
Robert	1
John	0
Maggie	0
Alex	1
Debbie	1

DP empowers the execution of functions on the database and the scrutiny of results. A pivotal question arises: "If someone is removed from the database, does the output of my function

change?" If the answer is no for each individual, the query preserves privacy perfectly, although its utility may be limited.

In contrast, without DP, even if modelers don't access raw data, malicious actors might reverse engineer model outputs, potentially revealing personal identities and sensitive information (just like in the Netflix case mentioned above). DP mitigates this risk by introducing random noise to the data, obscuring individual data points while preserving overall dataset properties. Modelers, aware of the randomness type, can still construct accurate group-level insights, while those attempting to exploit stolen data lack certainty about individual records.

To better understand how differential privacy (DP) functions, consider a hypothetical study investigating human inclinations toward sensitive topics, such as stealing. Traditional surveys that directly ask individuals whether they would steal something from an unattended shop if they could get away with it face significant challenges. Participants may hesitate to disclose controversial information, fearing judgment or potential consequences. If someone were to admit to stealing outright, they might be apprehensive about the information being leaked, making it challenging to obtain truthful responses.

DP addresses this issue through a method known as randomized response. In this approach, participants are given a degree of privacy, allowing them to respond truthfully while maintaining plausible deniability (the ability to deny any involvement in illegal or unethical activities because there is no clear evidence to prove involvement). Here's how it works:

Participants privately flip a coin twice. If the first coin flip results in a head, they answer the sensitive question honestly. However, if it lands on tails, they answer based on the second coin flip—responding with either a yes or no. This introduces an element of randomness. DP adds noise to responses, providing plausible deniability for individuals while allowing researchers to estimate the true distribution because the method preserves the overall dataset properties.

The core idea of DP lies in strategically adding the least amount of noise to data and queries to yield accurate results with optimal privacy protection (see also the section “Differential Privacy and the Privacy Gradient” in our article “The Most Common Data Anonymization Techniques”). There are two main use cases. The first one is local differential privacy, where a company needs to gather sensitive data points from distributed users. In this case, noise is added before sending data to the statistician (as illustrated in our example). The second is global differential privacy, where the data is already centralized in the company's databases, and the struggle is giving access to it either to internal users or external third parties. This case involves adding noise to the query outcome before its release from the database. Differential privacy's secret weapon is the art of adding noise strategically, granting plausible deniability to individuals in a database or training datasets.

PPTs in Practice, Case Studies

PPTs are increasingly gaining traction in real-world applications, with a prominent example in several areas. For example, a cross-pharma federated learning platform.[6] A consortium of

[6] Melloddy.eu

life sciences companies utilizes federated learning for collaborative drug discovery. By providing shared access to diverse datasets, participants can collectively enhance the predictive performance of drug discovery models, ultimately aiding in the identification of compounds for development. The platform is built on a centralized architecture incorporating machine learning algorithms and a robust privacy management system, ensuring secure and confidential data sharing among the participating entities.

In the realm of demographics, the latest US census has adopted differential privacy to safeguard individuals from identification risks while providing aggregated population data.[7] Additionally, the UN PETS (privacy-enhancing technologies) lab is actively testing various privacy-preserving technologies to foster collaboration among national statistics offices, researchers, and companies in handling shared data.

PPTs have demonstrated their effectiveness across diverse sectors in real-world applications. In healthcare, for example, homomorphic encryption enables secure collaboration between researchers and healthcare providers without compromising patient privacy. Financial institutions leverage these technologies to analyze transactional data securely, while e-commerce platforms employ them to enhance user personalization without exposing individual buying habits.

As these examples illustrate, PPTs are making a significant impact across different domains, ensuring the secure handling of sensitive data while fostering collaboration and innovation.

The impact of PPTs

Advantages and Disadvantages

PPTs have a profound impact on data protection across various dimensions, offering several advantages:

1. **Enhanced Security**: PPTs fortify security by safeguarding data during storage, processing, and transmission, addressing vulnerabilities at crucial points. This heightened security not only protects individuals but also fosters trust in businesses and institutions.

2. **Mitigation of Privacy Risks**: PPTs enable data analysis without exposing raw information, mitigating privacy risks. Organizations can extract valuable insights without compromising the confidentiality of individual data, aligning with user-centric data control principles.

3. **Regulatory Compliance**: PPTs provide a valuable tool for businesses navigating complex data protection laws. Features like differential privacy enable organizations to meet regulatory compliance requirements without compromising the quality of analytics and research.

However, the adoption of PPTs comes with certain trade-offs and challenges:

[7] https://www2.census.gov/library/publications/decennial/2020/census-briefs/c2020br-03.pdf

1. **Operational Delays**: when modelers lack direct access to data, they need to send models back to the data owner for execution, potentially slowing down the process of analysis.

2. **Computational Intensity**: techniques like homomorphic encryption, while effective, can be computationally intensive, impacting performance. Additionally, obscuring data with differential privacy may lead to reduced accuracy in specific use cases.

3. **Data Quality**: successful implementation of PPTs relies on good underlying data. Data owners must adhere to robust data management practices to handle anonymous queries effectively, especially since some modelers may not have direct access to the data.

4. **Privacy-First Approach**: PPTs should not be treated as mere add-ons but rather as fundamental components of the design process. Any process involving the sharing of private data should adopt a privacy-first approach to ensure effective integration and protection.

Future Trends and Developments

As technology evolves, so do privacy-preserving solutions. The integration of these technologies with emerging trends like AI and the Internet of Things presents a promising future. Machine learning models trained on privacy-preserving data and secure IoT ecosystems are on the horizon, opening new frontiers for innovation while maintaining a robust defense against privacy threats.

Businesses, government bodies, technology developers, and consumers all play pivotal roles in shaping the future of PPTs. From implementing these solutions responsibly to advocating for comprehensive regulatory frameworks, each stakeholder has a part to play in fostering a privacy-conscious digital ecosystem.

Conclusion

PPTs are not just safeguards against data breaches; they represent a paradigm shift towards a more ethical and secure digital future. However, no single technique serves as a silver bullet for privacy preservation. Achieving robust privacy requires the strategic layering of various technologies and thoughtful consideration of the right balance for each specific use case.

While PPTs offer substantial advantages in enhancing data privacy and security, their adoption necessitates careful consideration of trade-offs and a holistic approach to integration within the overall design of processes involving private data. By embracing these innovations, we can navigate the complexities of data protection with confidence, ensuring that privacy remains a fundamental right in the ever-evolving landscape of technology.

Real-time Analytics with PVML

PVML allows analysts to dive into real-time, online analytics without the worry of compromising privacy. This is done with mathematical models and differential privacy.

Understanding the challenges of existing PPTs, PVML has focused on creating a solution that does not require admins to alter the underlying data or infrastructure and does not require end-users to alter their existing workflow or query language. PVML's privacy engine can translate any query to a differentially private computation that can be applied on the live database. Learn more

This article was published on […] January 2023