# Privacy and Security Implications of Cloud-Based AI Services : A Survey

ALKA LUQMAN, Nanyang Technological University

RIYA MAHESH, Indian Institute of Technology

ANUPAM CHATTOPADHYAY, Nanyang Technological University

This paper details the privacy and security landscape in today's cloud ecosystem and identifies that there is a gap in addressing the risks introduced by machine learning models. As machine learning algorithms continue to evolve and find applications across diverse domains, the need to categorize and quantify privacy and security risks becomes increasingly critical. With the emerging trend of AI-as-a-Service (AIaaS), machine learned AI models (or ML models) are deployed on the cloud by model providers and used by model consumers. We first survey the AIaaS landscape to document the various kinds of liabilities that ML models, especially Deep Neural Networks pose and then introduce a taxonomy to bridge this gap by holistically examining the risks that creators and consumers of ML models are exposed to and their known defences till date. Such a structured approach will be beneficial for ML model providers to create robust solutions. Likewise, ML model consumers will find it valuable to evaluate such solutions and understand the implications of their engagement with such services. The proposed taxonomies provide a foundational basis for solutions in private, secure and robust ML, paving the way for more transparent and resilient AI systems.

## 1 INTRODUCTION

Machine learning (ML) is becoming increasingly ubiquitous nowadays, forming part of the computation fabric of many businesses. According to a Fortune Business Insights report [1], the global ML market size is expected to increase from USD 26.03 billion in 2023 to USD 225.91 billion in 2030. Various prominent cloud providers that used to offer infrastructure as a service (IaaS), platform as a service (PaaS) and software as a service (SaaS) now advertise Machine Learning as a service (MLaaS) in their portfolio. The cloud-based delivery model of ML as a service has reduced the barrier to entry in designing applications that use machine learning techniques. MLaaS offers machine learning algorithms, models, and tools that users can leverage without the need to build and maintain their own machine learning infrastructure. AI as a Service (AIaaS) is a broader offering that encompasses rule-based systems, symbolic reasoning, natural language processing APIs, chatbot builders etc., in addition to ML tools.

---

[1]https://www.fortunebusinessinsights.com/machine-learning-market-102226

---

### 1.1 Background

ML models, especially deep learning models which are becoming increasingly common, are susceptible to various attacks and pose serious safety implications in applications like smart homes, autonomous driving, etc. We show in Figure 1 how data poisoning attacks on a self-driving vehicle can introduce misleading data into the training dataset, causing the model to learn incorrect patterns. For example, attackers could manipulate the labels of camera images to mislead the model's understanding of the environment. Man-in-the-middle attacks can also exploit vulnerabilities in the model by performing evasion attacks. For instance, adversarial perturbations on road signs may cause the autonomous vehicle to misinterpret critical information, leading to potentially dangerous decisions. Robustness against these attacks is crucial to ensure the safety and reliability of autonomous driving systems, necessitating the development of resilient models and security measures.

With the proliferation of custom Large Language Models (LLMs) like ChatGPT and their usage, it is important to assess the privacy risks they present. This recent technology offering has been shown to be vulnerable to prompt injection attacks that expose the training instructions and files used to customize the custom chatbots [1]. Applications leveraging LLMs to read and parse personal emails to act upon their content can pose a serious security concern. Adversaries can craft emails containing specific prompts, forcing the LLM to execute malicious actions, like unauthorized purchases or content summarization with potentially harmful consequences.

### 1.2 Related Surveys

There has been a lot of research [2–4] focused on understanding the privacy and security risks associated with ML models. Various surveys have shed light on the vulnerabilities, potential attacks, and privacy concerns in the deployment of ML models. A taxonomy for privacy in ML proposed in [5] decouples using ML to perform targeted attacks and using ML as an aid to enhance user privacy. Attacks and their impact are judged in [6] based on the time of attack during its lifecycle.
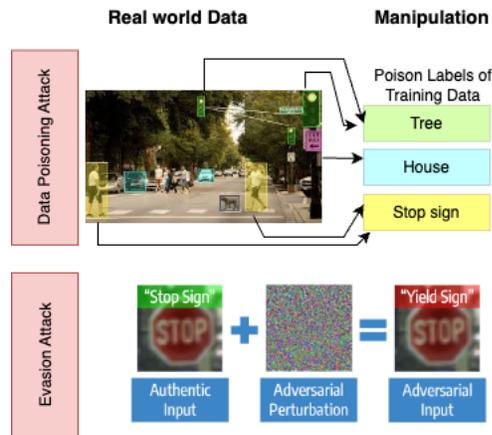


Fig. 1. Examples of attacks affecting autonomous vehicles during model training by poisoning input data labels or during inference by manipulating scenes in real-time.

These works have contributed significantly to our understanding of the subject. However, as the field of ML advances, the threat landscape evolves, demanding continuous exploration of novel challenges and innovative solutions. To address the ever-changing landscape of ML model security, there is a need for fresh perspectives that can provide insights into emerging threats and privacy-preserving techniques to combat them.

### 1.3 Motivation

The motivation of this paper is to highlight the importance of including privacy risks posed by ML models while promising data privacy guarantees in cloud services. Even if the training data is kept encrypted, the Deep Neural Network (DNN) models trained on such private data are still vulnerable to privacy leakage [7–9]. It is no longer enough to talk about data protection in silos, but to view it holistically with the ML models that make use of that data. ML attacks targeting edge IoT systems fall outside the scope of this paper. Here we focus on ML solutions hosted in the cloud via AIaaS.

### 1.4 Contribution

We propose a taxonomy to comprehensively view ML models and its training data based on resource ownership and usage rights. We catalogue the privacy guarantees assured by major cloud service providers like Microsoft, Google, Amazon, Huawei, Alibaba, IBM and Oracle, and map it to the proposed taxonomy. We include common attacks and their defence strategies presented in both theoretical settings as well as open source and private solutions.

The rest of the paper is structured as follows. Section II introduces the objectives commonly aspired for in information security and the many ways to achieve them - using protocols or specialized hardware. It then lays out the techniques used in the current cloud infrastructure for ML as a service to guarantee information security. Section III talks about the gaps in ML model security by listing out recent attacks. Section IV proposes the taxonomy considering ownership and usage rights and Section V contains the taxonomy considering defence techniques.

## 2 PRIVACY GUARANTEES AND THREATS IN THE CURRENT CLOUD ECOSYSTEM

### 2.1 Security Objectives

Information security encompasses a broad range of practices to protect sensitive information and prevent illegal access and modification of digital information. Many strategies are put forward to ensure the primary security measures of Confidentiality, Integrity and Availability, called the CIA triad. The CIA triad, depicted in Figure 2, serves as a foundational framework for designing and evaluating security measures and controls, along with other measures like non-repudiation and authentication.

**Confidentiality** is to prevent illegal access to the content of messages, ensuring that private and sensitive information is only accessible to authorized individuals. This is often achieved through measures such as encryption and access controls. Symmetric encryption techniques like DES and AES use a single public key to encrypt and decrypt data, while asymmetric encryption techniques like RSA and ECC use a pair of keys, an openly shared public key for encryption and a private key for decryption.

**Integrity** is to detect and prevent unauthorized modification of messages, ensuring that data and systems are reliable and trustworthy. This is usually achieved using cryptographic hash functions, Message Authentication Codes (MACs) and digital signatures. Hash functions generate unique checksums for data, and obtaining a different checksum indicates
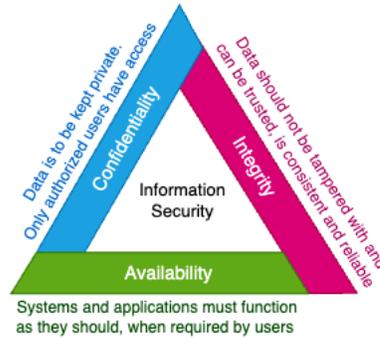
Fig. 2. CIA triad for Information Security.

that the underlying data has been changed. MACs are generated using a keyed hash function. Digital signatures provide a way to verify the authenticity of data and confirm that it has not been altered by unauthorized parties.

**Availability** is to safeguard against attacks that target the system, ensuring that the platforms through which data is accessed are accessible and operational when needed, without interruption. This involves safeguarding against downtime, denial-of-service (DoS) attacks, and other disruptions. There is no cryptographic measure to prevent this and it is often achieved using backups and redundancy strategies.

**Non-repudiation** is to prevent individuals from denying the validity or origin of actions or transactions. Digital signatures and audit trails help establish this.

**Authentication** is to verify the identity of users attempting to access information or systems. This is often achieved using public key constructs like digital signatures or private key constructs like Message Authentication Codes.

### 2.2 Cryptographic Standards

Cryptographic standards refer to established ciphers for various applications, achieving security objectives mentioned in the previous subsection. Some of the commonly used algorithms are mentioned below.

**Private-Key Cryptography**

- **AES (Advanced Encryption Standard)** is a symmetric-key algorithm that operates on data blocks of size 128 bits with keys of length 128, 192, or 256 bits. It was developed in a public competition organized by the U.S. National Institute of Standards and Technology (NIST), where the Rijndael algorithm developed by Belgian cryptographers Vincent Rijmen and Joan Daemen, was selected as the winner. AES encryption involves multiple rounds of transformations on the data block, including substitution, permutation, and mixing operations. These rounds are performed on the data using a round key derived from the original encryption key. AES, with its larger key size, is considered more secure than its predecessor DES. It is usually implemented using block ciphers for encryption. The two popular cipher modes are *CBC (Cipher Block Chaining mode)* and *GCM (Galois/Counter mode)*. CBC encrypts data block using XOR chaining while GCM operates with counter mode of encryption and Galois mode for authentication leading to a higher throughput.
- Lightweight block cipher algorithms like PRESENT, SIMON and SPECK are designed to provide secure encryption with a focus on simplicity and efficiency in resource-constrained environments. **PRESENT** is known for its small hardware footprint and is used in low-power devices like RFIDs and sensors. It uses a compact and

efficient round structure with a minimal number of operations, making it well-suited for resource-constrained devices. **SIMON** is designed to offer strong security while being highly efficient in terms of both hardware and software implementations. It can be used in a wide range of applications, including lightweight cryptography for small IoT devices, secure communication protocols, and encryption in constrained computing environments. **SPECK** is a family of lightweight block ciphers designed to provide efficient and secure encryption for various applications such as IoT devices, mobile communication, and constrained computing platforms.

**Public-Key Cryptography**

- **RSA (Rivest-Shamir-Adleman)** is an asymmetric-key algorithm which uses a public key for encryption and a private key for decryption. When a user wants to send an encrypted message to a recipient, they use the recipient's public key to encrypt the message. This ensures that only the recipient, who possesses the corresponding private key, can decrypt and read the message. The operations in RSA involve modular exponentiation with large prime numbers, making it computationally infeasible to derive the private key from the public key.
- **ECC (Elliptic Curve Cryptography)** is a type of public-key cryptography that is based on the algebraic structure of elliptic curves over finite fields. The strength of ECC lies in the difficulty of calculating the private key from the public key, known as the elliptic curve discrete logarithm problem. It is widely used for securing digital communications, data encryption, and digital signatures due to its strong security properties and efficiency. ECC keys are much shorter than equivalent RSA keys, giving it an advantage in resource-constrained environments.
- **PQC (Post-Quantum Cryptography)** refers to the field of cryptography that focuses on developing cryptographic algorithms and protocols that remain secure against attacks from quantum computers. Quantum computers have the potential to break many of the widely used encryption schemes like RSA and ECC, by leveraging quantum algorithms like Shor's algorithm for factoring large numbers and Grover's algorithm for searching unsorted databases.

**Message Authentication Code and Hash Functions**

- **MAC (Message Authentication Code)** is a cryptographic technique used to verify the authenticity and integrity of a transmitted message. It is derived from the message and a secret key, and is sent alongside the message to ensure that it has not been tampered with during transmission. It utilizes a secret key, which is private to the sender and the recipient, to both generate and verify the MAC.
- **MD5 (Message Digest) Algorithm** is a cryptographic hash algorithm that generates a hash code of fixed length (128 bits) for any arbitrary input message. While still in use, the algorithm has shown vulnerability to collision attacks and is found to be cryptographically broken. This has resulted in the adoption of alternative algorithms like SHA for hashing.
- **SHA (Secure Hash Algorithm)** converts any input into a condensed hash value of fixed length called as the Message Digest. It is based on Merkle-Damgard construction technique. *SHA-0*, *SHA-1*, *SHA-224*, *SHA-256*, *SHA-384*, *SHA-512* are the existing versions of this family which differ in their strength of security and size of the block that they operate upon.

### 2.3 Communication Protocols

Cryptographic communication protocols are sets of rules and procedures that govern secure communication between two or more parties over a network or any communication channel. These protocols employ cryptographic techniques to ensure the objectives of confidentiality, integrity, and authenticity of the data being exchanged. Some examples are given below, applicable from lowest to highest layers of the OSI stack.

- **SSL (Secure Sockets Layer)** is a cryptographic protocol designed to provide secure communication over a computer network. It was first developed by Netscape Communications in the 1990s to add the HTTPS protocol to their web browser Netscape Navigator. SSL was designed to encrypt data transmitted between a client and a server, ensuring that the data exchanged during these sessions could not be easily intercepted or deciphered by malicious actors. Due to various security vulnerabilities, SSL has been deprecated in favour of TLS.
- **MACsec** is a data-link layer encryption method using the IEEE 802.1AE MAC Security Standards. Packets are encrypted and decrypted on the devices before being sent, preventing snooping or wiretapping attacks. This technology is integrated on the network hardware, and so provides encryption on the network hardware with no measurable increase in link latency.
- **TLS (Transport Layer Security)** is a cryptographic protocol designed to secure communications over a computer network. It is commonly used to protect data transmitted over the internet, ensuring the privacy, integrity, and authenticity of the data exchanged between two systems. It is commonly used in securing web browsing (HTTPS), email communication (SMTPS, IMAPS), VPN connections etc. TLS 1.3, defined in August 2018, is the most recent version.
- **SSH (Secure Shell)** is a cryptographic network protocol used for securely connecting to and managing network devices and servers over an unsecured network. SSH was first designed in 1995 and uses public-key cryptography to allow users to authenticate themselves to a remote computer. It also includes utilities like SCP (Secure Copy Protocol) and SFTP (SSH File Transfer Protocol) for secure file transfers between systems, with encryption and data integrity checks. OpenSSH is the most widely implemented variant.
- **S/MIME (Secure/ Multipurpose Internet Mail Extensions)** is a standard for securing email messages with encryption and digital signatures. It is commonly employed in email clients and servers to protect the confidentiality of email content (via encryption) and to verify the authenticity of the sender along with non-repudiation of origin (using digital signatures).
- **VPN (Virtual Private Network)** is a protocol that provides secure communication over the internet, allowing users to transmit data privately and securely. It is often used to create secure, encrypted tunnels between clients and servers. It is widely used by organizations to ensure confidentiality of the data transmitted over the network.

### 2.4 Trusted Computing

- **HSM (Hardware Security Modules)** are robust computing devices that provide cryptographic security to data by performing encryption, decryption and key management processes. They are embedded on secure cryptoprocessor chips and assist in creating digital signatures and certificates.
- **Trusted Execution Environment** [10] is a secure enclave isolated from all system software or hardware outside it. It provides a trusted platform to carry out computational tasks and protects the integrity of model and data from an untrusted host environment during execution. TEEs play a crucial role in providing both

server-side and client-side protection[11]. The former focuses on preventing threats during model training and inference stages while the latter involves securing the model and client data from vulnerabilities once deployed for performing tasks.

- **Homomorphic Encryption** [12] is a cryptographic technique that allows computation to be performed on encrypted data without having to decrypt it. Mathematical operations like addition, multiplication etc. can be performed on the ciphertext directly through specialized encryption algorithms that maintain certain mathematical relationships between plaintexts and ciphertexts. One popular homomorphic encryption-based MPC protocol is the BFV protocol, which allows parties to securely compute a variety of functions on their private inputs.

- **Secure Multi-party Computation** is a technique whereby multiple parties collaborate to compute results on their private data, without revealing each part. This can be used to perform privacy-preserving data analysis, privacy-preserving machine learning, enabling smart contracts in blockchain, etc. MPC protocols are based on Secret sharing, Garbled circuits, Threshold cryptography-based MPC etc.

- **Searchable Symmetric Encryption** [13] is a cryptographic technique that allows search on encrypted records in databases. The advantage of SSE is in the ability to selectively retrieve parts of documents. SSE is well-suited for scenarios where users need to search within a cloud-hosted encrypted database without revealing the plaintext data. Most SSE schemes today utilize secure indexes generated from pre-processed document-keyword pairs while encrypting documents using a symmetric encryption algorithm. The server compares query tokens with messages in the index and matched documents are returned to the data user.

### 2.5 Current privacy guarantees

Data outsourcing is a widely practised cloud service where users and companies enlist third parties to store and manage their data. This is favourable in the face of high costs of storage and of skilled professionals to design and manage secure in-house storage systems. But this comes at the cost of potentially exposing private data [14]. Encryption is what enables data privacy on the Internet today. If an entity Alice is sending a private message to Bob, as shown in Figure 3, the message is encrypted and the only parties with access to the key are Alice and Bob. The private data cannot be understood by any third party who tries to intercept the message while it is in transit through a communication channel.

Existing cloud providers provide data encryption at storage nodes (Data at rest), communication links (Data in transit) and compute nodes (Data in use). This categorization is according to the 3 states of digital data.

(1) **Data at rest** refers to data that is stored on physical or virtual storage devices. Data encryption at rest involves encrypting this data to protect it from unauthorized access if the physical storage media is compromised or stolen. It requires customers to use effective key management systems. The keys are stored securely and rotated regularly to minimize any risk of exposure.

(2) **Data in transit** refers to data that is being transmitted between devices or services within a cloud environment or between the cloud and external systems. Data is encrypted before it leaves the source and decrypted upon arrival at the destination to maintain confidentiality. Common protocols for encrypting data in transit include HTTPS for web traffic, SSL/TLS for secure communication, and VPNs for secure network connections.

(3) **Data in use** refers to data that is actively being processed by applications or services in the cloud. This data resides in computer RAM, CPU caches, registers etc. Cryptographic protocols like homomorphic encryption
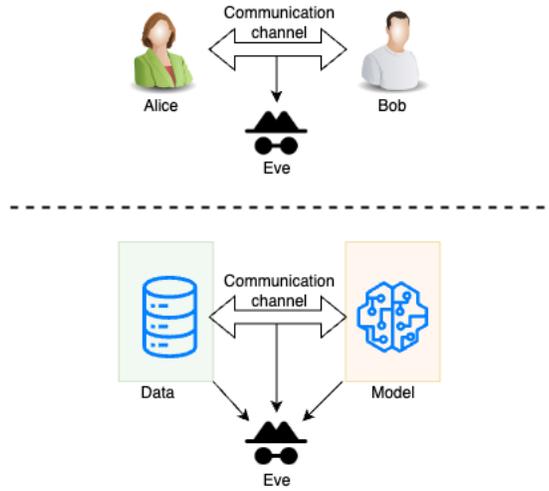
Fig. 3. Extending the threat model of cryptography systems to the ML ecosystem.

and secure multi-party computation perform computation operations on encrypted data. Another alternative is the use of secure enclaves, which maintains the data in encrypted form in RAM, but in clear text inside CPU caches and registers.

The underlying technologies at each step are consolidated in Tables 1 and 2 to see how privacy is currently guaranteed in some major cloud providers, for data and ML models respectively. Confidential computing is overwhelmingly used to address attacks on data during processing by isolating it exclusively to the CPU handling the processing. The information in the tables also emphasize how the guarantees for model privacy currently are piggybacking on existing definitions of data privacy. In light of the widespread adoption of ML applications, these distinctions are no longer enough to encompass all possible interplays of ML model usage with the privacy of its underlying training data. There is a wide range of attacks on ML models which expose additional attack surfaces that a malicious actor may exploit to compromise the privacy of personal and sensitive information. The training data becomes an additional attack vector that threatens the robustness of the ML model.

The artefacts in a cloud based ML solution include

(1) Training dataset - This includes the private images or textual information used to train an ML model.
(2) Trained ML model - This can be the structure of the tree and the decision criteria in case of a decision tree, the weights of a DNN, or a reward matrix in RL.
(3) Test dataset - It is used to statically evaluate the model performance.
(4) Inference data - This is the private data of consumers who use the service.

In the subsequent discussions of the taxonomy we use the term *DATA* to refer to the training dataset, test dataset and inference data, while *MODEL* refers to the trained ML model. An *ML model consumer* is an entity or system that utilizes the model for predictions or decision-making and *ML model provider* is the entity responsible for creating, maintaining, and supplying the ML model and its related services to the consumers. Table 3 categorizes the numerous attacks on

Table 1. Data privacy guarantee in current cloud ecosystem.

| Cloud provider | At rest | In transit | In use |
|---|---|---|---|
| MS Azure [15–17] | AES-256 in CBC mode<br>FIPS 140-2 compliant | VPN (utilizing IPsec/IKE encryption)<br>TLS 1.2 & 1.3<br>SSH (on Linux VMs)<br>SMB 3.0 (between Azure virtual networks)<br>IPSec (between on-prem network and Azure network)<br>MACsec (LAN between datacentres) | Confidential computing<br>(Intel SGX) |
| AWS [18, 19] | AES-256 | TLS 1.2, AES-256<br>AWS Key Management Service, hardware based HSM | Confidential computing<br>(Nitro enclave) |
| GCP [20] | Disks created after 2015 use AES-256 (using library Tink)<br>Disks created before 2015 use AES-128<br>FIPS 140-2 compliant | TLS (default) uses Google's implementation called BoringSSL<br>IPSec tunnels<br>S/MIME for Gmail<br>Managed SSL certificates<br>AES-128 (ALTS) for Layer 7 encryption | Confidential computing (AMD Secure Encrypted Virtualization) |
| Huawei [21, 22] | AES-256<br><br>Key Mangement Service | TLS 1.2, SSL, VPN | Trusted Intelligence Computing Service<br>(ARM Trusted Firmware v1.3) |
| IBM [23] | AES - 256 in CBC mode<br>RC4-128 encryption with MD5-128<br>AES-128 encryption with MD5-128<br>AES-256 encryption with SHA-256 | TLS 1.3<br>AES in CBS or GCM mode<br>RSA, ECDSA, ECDHE algorithms | Confidential Compute [24]<br>(Intel SGX)<br>Fortanix Runtime Encryption Platform |
| Oracle [25] | AES 256 SHA-256 with RSA<br>TDES (Triple Data Encryption Standard) 168 bits<br>Regional Encryption algorithms : ARIA, SEED and GOST | TLS 1.1 & 1.2 | Confidential computing<br>(AMD Infinity Guard) |
| Alibaba [26] | HSM, AES-256<br>KMS at 2 levels Customer Master Key and Data Encryption<br>Key | SSL, TLS 1.2 & 1.3<br>VPN Gateway<br><br>Alibaba Cloud Smart Access Gateway | Intel SGX enclaves |

Table 2. Model privacy guarantee in current cloud ecosystem.

| Cloud provider | At rest | In transit | In use |
|---|---|---|---|
| MS Azure [16] | Encrypt model using Azure managed encryption key | Encrypt network traffic using TLS 1.2 | Confidential computing |
| AWS [27] | Encrypt model using AWS managed encryption key | Encrypt network traffic using TLS 1.2 | Confidential computing |
| GCP [28] | Encrypt model using Customer Managed Key | Encrypt network traffic using TLS | Confidential computing |
| Huawei [21] | Encrypt model using Customer Managed Key | Encrypt network traffic using TLS 1.2 | Trusted intelligence computing service |
| IBM [24] | Encrypt model using Customer Managed Key | Encrypt network traffic using TLS 1.3 | Confidential computing |
| Oracle [25] | Encrypt model using Customer Managed Key | Encrypt network traffic using TLS | Confidential computing |
| Alibaba [26] | Encrypt model using 2 Customer Managed Keys | Encrypt network traffic using TLS | Confidential computing |

data and model which have been proposed in the recent past. The attacks mentioned in this paper are a representative selection of the works till date and not an exhaustive list.

Table 3. Attacks on Data and Model

| Attack type | Attack name | Description | Examples | Access Type |
|---|---|---|---|---|
| Exploratory | Adversarial Transfer Attack | Attacker trains a substitute model to study and transfer the attacks on this model to the target model | Transfer Attack[29] | Blackbox |
| | Membership Inference Attack | Adversary determines if a particular data point was a member of the training dataset | Shadow Attack[30] | Whitebox |
| | | | GAN based Membership Inference attack[31] | Blackbox |
| | Model Inversion Attack | Attacker aims to reconstruct the training data by accessing the model | Activation Maximization [32] | Whitebox |
| | | | GAN and Reinforcement Learning based Model Inversion[33] | Blackbox |
| | Model Reconstruction Attack | Adversary reconstructs the trained model | Local Model Reconstruction Attack[34] | Whitebox |
| | Attribute Inference Attack | Adversary reconstructs the model attributes from its local update stages | Attribute Reconstruction Attack [35] | |
| | | | Model Inversion Attribute Inference Attack [36] | Blackbox |
| | Label Inference Attack | Attacker aims to infer the label of the data points | Label Inference Attack against Vertical Federated Learning[37] | |
| | | | Label Only Membership Inference Attack[38] | Blackbox |
| | Infer Class Representatives | Attacker infers the features of every class to construct representative samples for each of them | Generative Adversarial Network-based infer class representative attack[39] | Blackbox |
| Evasion | Model Evasion | Crafting adversarial examples to introduce perturbations to the input data based on optimization techniques to misclassify them | Carlini and Wagner attack[40] | Whitebox |
| | | | Zeroth Order Optimization[40] Deepfool[40] Fast Gradient Sign Method[40] | |
| | | The attacker uses trained Generative Adversarial Network to introduce adversarial features to evade the model | GAN Attack against intrusion detection [41] | Blackbox |
| Poisoning | Data Poisoning | Injecting adversarial action choices into the agent's trajectory | Action poisoning Attack[42] | Blackbox |
| | | Attacker crafts inputs with similar feature representation provoking the model to misinterpret the data | Feature Collision Attack[43] | |
| | | Adversary mislabels the input dataset | Label Flipping Attack[44] | |
| | | Attacker attaches a patch to the image that leads to misclassification of the data into an incorrect class | Adversarial Patch Attack[45] | Whitebox |
| | Model Poisoning | Poisoning local model updates during the training phase | Whitebox online model poisoning attack on Production Federated Learning[46] | Whitebox |
| | | Adversary poisons the model architecture or data during training phase with hidden triggers which are exploited during the deployment stage to misclassify data | Backdoor Attack on CNN [47] | Whitebox |

## 2.6 Threat model in the Cloud

A knowledge base of real-world attacks is provided by MITRE ATT&CK, to be used by organizations who want to develop effective cybersecurity practices. They maintain a continuously evolving list of attack techniques [2] which grants an attacker different levels of access to a cloud enterprise solution that provides IaaS or SaaS. Attacks are grouped by reason of attack a.k.a tactics, i.e. the goal of the adversary during each action. The different goals of an adversary trying to attack at various stages of an enterprise solution is listed in Table 4. The matrix also provides different tried and tested approaches for the attack's detection and mitigation.

Table 4. Tactics listed in the ATT&CK matrix curated by MITRE for Cloud Enterprise. Highlighted rows represent tactics which are also applicable to modern ML solutions.

| Tactic | Explanation |
| --- | --- |
| Initial Access | Techniques used to gain an initial foothold within a network. |
| Execution | Techniques by which the adversary runs malicious code on a local or remote system. |
| Persistence | Tactics that the adversary uses to maintain their foothold across any interruptions like system restarts or credential changes. |
| Privilege Escalation | Tactics by which the adversary tries to gain elevated permissions on a system or network like root level or administrator access. |
| Defense Evasion | Methods that the adversary uses to evade detection while the system is compromised. |
| Credential Access | Methods used by the adversary to steal user login credentials. |
| Discovery | Techniques used by the adversary to explore and gain knowledge about the system or network they are attacking. |
| Lateral Movement | Techniques the adversary uses to move through the system once it has entered, to reach the target of the attack. |
| Collection | Techniques used by the adversary to gather information about the system which will aid in the attack. |
| Exfiltration | Techniques to steal the data. |
| Impact | Techniques to disrupt the availability of a system or destroy its data. |

Certain tactics in this categorization also apply to ML solutions. In the *persistence* tactic, an adversary maintains prolonged access to the system. This is manifested in Backdoor ML attacks, where malicious actors manipulate an ML model during its training or deployment to introduce a hidden behavior, similar to a backdoor or trojan. This backdoor behavior is designed to trigger specific responses or outputs when the model encounters specific, often subtle, input patterns or trigger conditions. Such attacks can compromise the integrity and security of the ML system. Attacks to infer the presence of certain samples in the data used to train an ML model, including its statistical properties demonstrate the tactics of *collection* and *exfiltration*. Evasion attacks in ML refer to the process of crafting perturbed input data, called adversarial examples [48, 49], to misclassify it during inference or testing. Such adversarial techniques along with attacks that manipulate the model inference at run time exhibit the *impact* tactic. We review these attacks in detail in the next section, after a brief primer on ML techniques.

---

[2]https://attack.mitre.org/matrices/enterprise/cloud/

## 3 THREATS IN THE ML ECOSYSTEM

### 3.1 Machine Learning Overview

Machine Learning (ML) has witnessed remarkable growth in the past decade, transforming how we interact with technology, handle data, and solve complex problems. There are three fundamental paradigms in **Classical ML**, based on the nature of the feedback available to the learning system.

- **Supervised Learning** trains a model on labeled data with explicit input-output pairs. Example use cases include image classification, stock price prediction via regression etc.
- **Unsupervised Learning** discovers patterns or structures in data without explicit labels. Some applications that use this include customer segmentation via clustering, fraud detection and generative modeling.
- **Reinforcement Learning** trains intelligent agents to make a sequence of decisions in an environment in order to achieve a goal or maximize a cumulative reward. It involves exploring the environment through interaction with it and exploiting the information present in it. It is commonly used in urban traffic management, mastering motor control in robotics and playing strategy games like Go, Poker, Dota2, etc.

**Deep Neural Networks (DNNs)** introduced the concept of neural networks with multiple hidden layers. Convolutional Neural Networks (CNNs) are well suited for working with image data on tasks like image recognition and image segmentation. Recurrent Neural Networks (RNNs) perform well with sequential data, and are used in speech recognition and language translation. **Transformer Networks** use attention mechanisms [50] to process sequences of data in parallel, enabling unprecedented improvements in natural language processing and machine translation. They are also used in computer vision and reinforcement learning.

**Transfer Learning** [51] is a technique that extends the knowledge learned on one domain to a different but related task in another similar domain. Transfer learning is especially valuable when there is limited data to train a target task, as it leverages the knowledge acquired from a larger or more relevant source task. Models like BERT highlight the power of pre-training and fine-tuning required for such tasks.

**Federated Learning** [52] is a collaborative learning approach that allows multiple devices or entities to collaboratively train a shared model while keeping their data decentralized and private. Instead of sending raw data to a central server for training, federated learning reduces privacy concerns by training models locally on each device and sending the local model parameters to a central server for aggregation.

**Split learning** [53] is another privacy-preserving collaborative learning approach that shares activation maps between devices after using Homomorphic Encryption to encrypt them. Activation maps represent the activation of neurons in a neural network generated during forward propagation in response to input data. They can provide insight into the features learnt at different layers of the neural network.

### 3.2 Privacy Attacks

Privacy Attacks refer to practices that aim to infer private information of participants, be it individuals or organizations. While data at rest is vulnerable to theft by malicious agents, deep learning models are also vulnerable to reveal their underlying training data and its characteristics due to deep leakage via gradients. In these attacks the private data used to train ML models and the model itself are the assets under attack.

**Data privacy attacks** refer to techniques where the adversary tries to leak private and sensitive information from the dataset used to train an ML model. Some examples of such attacks include -
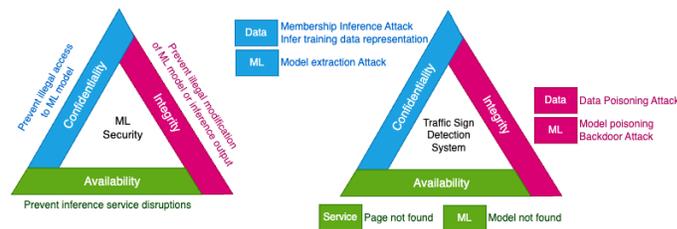
Fig. 4. Extending the CIA triad for Model Security and using it to depict the security attacks possible on an example ML application that identifies physical traffic signs on roads to control driving in autonomous vehicles.

- Inference of class representatives - Such attacks deduce the features of importance to different classes in the training data and generate samples in real time, which appear to come from the same distribution as the training data.
- Inference of statistical properties - These attacks attempt to infer satistical features like when did a particular sample first appear in the training set, the ratio of men to women etc.
- Inference of sensitive attributes - In these attacks the adversary leverages additional public datasets to re-identify individuals or their private information used to train an ML model. Given partial information from a patient's medical record, [54] attempts to infer each individual's genotype information.
- Reconstruction of training input and labels - These attacks target ML models to reverse engineer or recover sensitive information from the model's training data. The attacker typically has access to the model's predictions and tries to infer the inputs or labels that led to those predictions.
- Membership Inference Attacks - These attacks gauge if an input sample was part of the training dataset used to train the ML model. It is usually exposed by non zero gradients in the NN embedding layer, or by overfit or complex models with high number of parameters. This is a Black Box attack, which can succeed without direct access to the ML model, just by listening to communication channel. Mitigating membership inference attacks often involves adding noise to the training data or applying privacy-preserving techniques like differential privacy to protect against data leakage.

**Model privacy attacks** are methods by which an attacker tries to steal the ML model itself, in the case of models which are not freely available to the public. These include -

- Model inversion attacks try to infer training data using query-response pairs by observing the predicted label and confidence scores of prediction. This is usually a black box attack, but can become a White Box attack, after a successful extraction attack on the model.
- Model reconstruction attack is a superset of the model inversion attack. A GAN reconstruction is used to recreate the ML model. This attack is also applicable to non neural network ML models which store feature metadata.
- Model extraction is a black box attack that tries to learn the same task as the target with at-par accuracy, by learning where the decision boundaries of the target model lie. It has been shown that hyperparameters can also be successfully extracted.

### 3.3 Security / Robustness Attacks

Robustness Attacks aim to compromise the global ML model by manipulating NN layer gradients or poisoning the training data. In these attacks the quality of the ML model is the commodity of value.

**Data robustness attacks** target the training data residing on devices and attempts to modify it so that the ML model does not reach acceptable accuracy levels.

- Data poisoning attacks like the Label Flipping attack modify the target class of training data.

**Model robustness attacks** tamper with the model training process to degrade the prediction accuracy of individual or global models.

- Model poisoning attacks like the Byzantine attack work by devices uploading malicious gradients so that the averaged model moves further away from the global minima. Sybil attacks are another form of model poisoning where multiple clone devices are spawned and controlled by one powerful device to tamper with the gradient convergence.
- Backdoor attacks embed a backdoor task during model training, which gets activated by a specific input. Examples of such attacks include Adversarial Backdoor Embedding, Clean Label feature Collision etc.
- Evasion attacks are white box attacks that add small perturbations to the model input in order to change the predicted label.

### 3.4 Threat Model for an ML service

Threat models are used in cybersecurity to systematically identify and mitigate potential risks in a system. Resources which are targeted by an attacker are called assets. In any ML application, the most important assets are the ML model and the data used to train it. Entities that interact with the system or play specific roles are called actors. Here we focus on the data owner who provides training data, the model owner who builds and trains the ML model, and the model
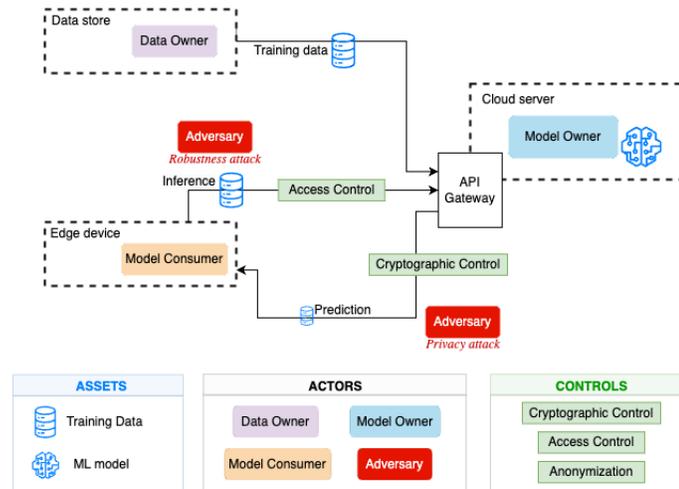


Fig. 5. Threat model on a Data Flow Diagram of a simple ML-as-a-service model. The actors interact with the system using assets, which are the resources targeted by an attacker. The dashed lines represent threat boundaries. Controls are placed at vulnerable attack surfaces to mitigate risk.

consumer who uses this ML service for inference. The other actor is an adversary or a threat agent, which can either be active or passive. A passive adversary is an honest-but-curious participant who tries to leak private information about the underlying training data. An active adversary is a malicious participant who tries to sabotage the ML model's performance.

Trust boundaries are drawn to logically separate areas where trust is assumed and areas where trust is not assumed. In the threat model shown in Figure 5 the dashed lines represent trust boundaries among different devices. Controls are countermeasures put in place to mitigate the vulnerabilities identified in the system. Preventions are controls that may completely prevent the possibility of a particular attack, e.g. access controls to authorize certain operations. Mitigations are controls to reduce either the likelihood or the impact of a threat, e.g. anonymization defences like differential privacy, cryptographic controls that encrypt data etc.

### 3.5 Threat Model considering the lifecycle of an ML model

Figure 6a depicts a threat model which takes into account the time when an attack occurs in the lifecycle of an ML model. This criterion provides a deeper insight into the precise phases of model design and deployment which exhibit vulnerability. This allows application developers to choose appropriate defence strategies at different stages.

The lifetime of an ML model comprises of the following stages:

- Data Collection - The first step under data preparation involves obtaining data in its raw form. Tampering the dataset at this stage will propagate the error throughout the subsequent phases.
- Data Cleaning - The data from the previous step is filtered by removing noises and detecting anomalies. Disregarding useful data as unwanted could mislead the model.
- Data Labeling - Assigning label to data points, usually in the context of supervised machine learning is the next step. Attacks manifest through intentional mislabelling or manipulating the data to provoke the model for erroneous classification.
- Feature Engineering - It involves selection, transformation and augmentation of data to extract relevant features and attributes as inputs to the model. Introducing false data that mimics the original features or inducing the model to generate adversarial outputs by strategically biasing specific features are the possible forms of attack vulnerabilities.
- Model Training and Evaluation - Carrying out attacks during the training and evaluation phase involves manipulating the learning trajectory of the model, modifying gradients and introducing perturbations to inputs within the hidden layers.
- Model Inference - Attacks primarily focus on exploiting the model after training to extract information about the model parameters, reconstruct the model and infer the membership and the representative class of a particular data point.

The attacks listed in Table 3 are mapped to these stages where applicable in Figure 6a and their corresponding defences are also mentioned.

## 4 TAXONOMY CONSIDERING OWNERSHIP AND USAGE RIGHTS

We will now merge the cloud security guarantees, as discussed in Section 2, with the threats presented Section 3 to provide a new perspective through which to analyze the cloud ecosystem.
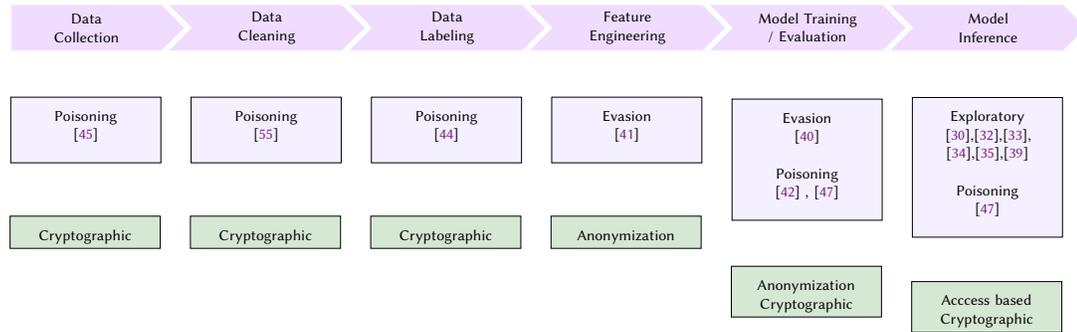
(a) The first row breaks down the ML lifecycle into stages. The attacks (detailed in Table 3) and defences applicable at each stage are shown in the second and third rows respectively.



(b) During model training



(c) During model inference
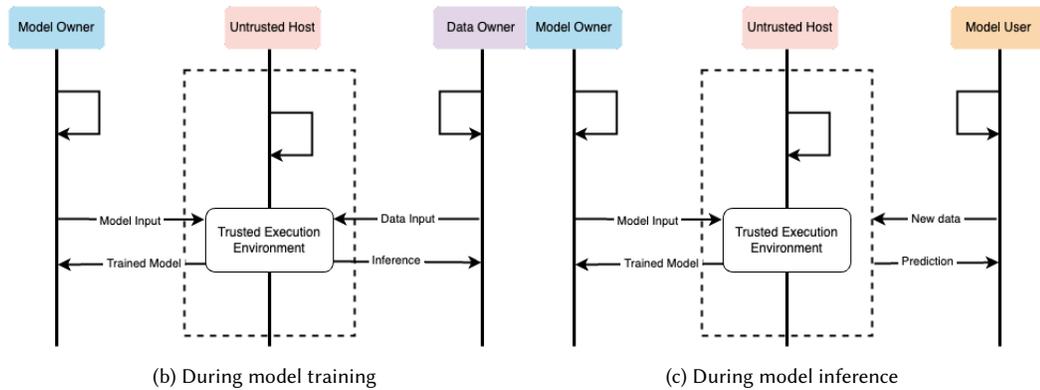
Fig. 6. Threat model on different stages during the lifecycle of an ML model is shown in Figure 6a. The attack surfaces involved during 2 specific stages is further elaborated in Figures 6b and 6c.

In today's world of connected devices, the data required to train an ML model may not all be present on one device. Data is generated by multiple autonomous devices in Wireless Sensor Networks [56], which work together in the Internet of Things [57] to achieve one or many objectives like traffic prediction [58], learning smart home policies [59] etc. With the advent of cloud computing, services like Google's BigTable [60] have demonstrated better flexibility and higher performance by distributing storage resources rather than having the data aggregated at a single data centre.

Legal restrictions like the General Data Protection Regulation (GDPR) [61] regulate how data of people in the European Union is collected, stored and processed. Similar legislations in other countries include the Personal Data Protection Act (PDPA) [62] in Singapore, the Data Protection Act in the United Kingdom [63] and the Consumer Privacy Bill of Rights (CPBR) [64] in the United States. These policies deter or complicate the process of aggregating data across devices and regions.

We propose a taxonomy inspired by Flynn's taxonomy which is used to classify computers based on the number of instruction streams and data streams in their architecture. Flynn's taxonomy was introduced in 1966 to classify parallelism in computers and it captures what *instruction* each processor executes along with what *data* it works on.

We take a similar view to classify ML applications based on *model objectives* and the distributed nature of the *data* used to train it. Figure 7 presents examples for each setting in the proposed taxonomy.

### 4.1 Single Data Single Model (SDSM)

Amongst the devices participating in ML, the training data can reside at a single device or be distributed across devices, with varying levels of trust between these devices. The data owner trains a single ML model which is then available for private or public use.

This includes applications where one ML model is trained on data present on the same machine or data sourced from a different machine, e.g. applications using third-party ML as a Service. It also includes ML applications that aggregate data from multiple machines that still fall under the purview of a single owner, e.g. inventory management and real-time vehicle tracking for logistics optimization, smart farming, remote sensing and diagnostics.

### 4.2 Single Data Multiple Model

A single data owner with access to a large corpus of data can create multiple ML models such that they still retain ownership of the models, but it is available for use by anyone. This includes Multi-task learning applications where commonalities and differences across tasks are leveraged to learn multiple ML objectives from a shared dataset. [65] use head pose estimation and facial attribute inference as auxiliary tasks for facial landmark detection; [66] jointly predict the phoneme duration and frequency profile for text-to-speech conversion. Applications that use transfer learning, like detecting Alzheimer's disease [67] using a modified AlexNet model architecture pre-trained on ImageNet, also fall into this category.

### 4.3 Multiple Data Single Model

In situations where a large public corpus of data is not available, we can aggregate the learning from multiple private datasets. Storing and processing large volumes of data using a cluster of machines is well-studied in distributed database systems and distributed computing. Tools like Apache Hadoop and Apache Spark used prevalently in Big Data analytics utilize paradigms like MapReduce to perform statistical aggregation and draw insights from multiple data sources.

The numerous devices or clients represent multiple data owners whose datasets may differ in volume and quality, and cannot be collated in a single place for training a central model. All the devices that contributed towards training a global model have the right to use it for their own predictions. Google's GBoard uses personal keystrokes from multiple mobile phones to train a Next Word Prediction model via Federated Learning. BrainTorrent [68] uses Federated Learning in the domain of medicine to perform the task of whole brain segmentation in a decentralized setting.

### 4.4 Multiple Data Multiple Model

This category encompasses solutions that leverage a globally learnt model via Federated Learning and further personalize it to suit different clients [69]. All participants own their own finetuned version of a global model for their individual use. Personalized Federated Learning can be achieved by treating the task like a meta-learning [70] or multi-task learning [71] problem or by mixing the global and local models.

Examining the ownership and usage rights of data within cloud provider systems is crucial for ensuring robustness in design. Understanding the spread and statistics of the data allows for better management of access and usage policies, better anomaly detection during the model training process, etc. Implementing distributed training policies to leverage
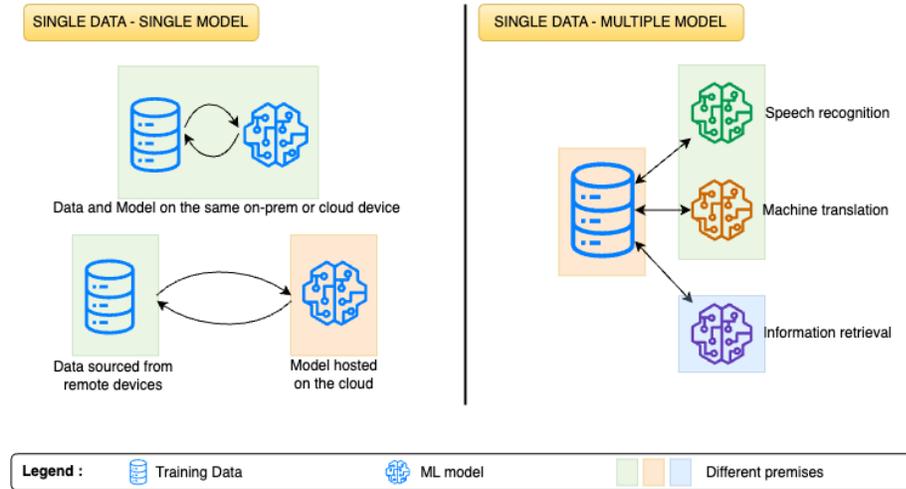
(a) Applications with single data owner
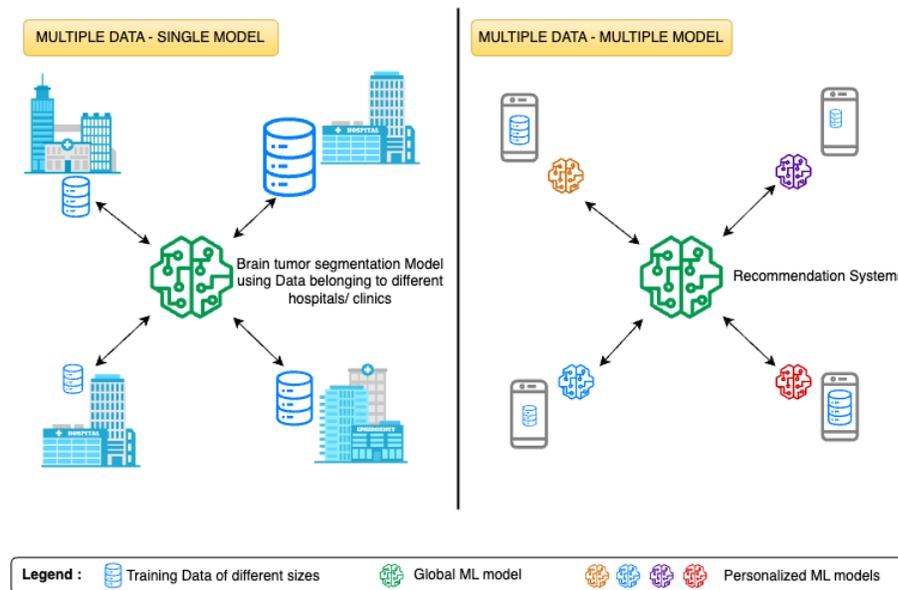


(b) Applications with multiple data owners

Fig. 7. Proposed taxonomy of Data-Model organization considering ownership and usage rights.

distributed data also needs to balance efficient communication strategies with effective learning strategies. This will help to create useful ML models that are amenable to a wider user base while still being technically robust.

## 5 TAXONOMY CONSIDERING DEFENCE TECHNIQUES

The need for different ML defences in the real world arises from the diverse range of security, privacy, fairness, and operational challenges that ML systems face. These defences are essential to ensure the responsible and secure deployment of ML models in critical applications and to maintain user trust and confidence in AI technologies.

A categorization of attack types that ML models face is presented in Table 3. The defense strategies employed to mitigate these specific attacks can be classified into 3 buckets.

(1) **Cryptography based** defences guarantee confidentiality and integrity of data.
   - Symmetric key cryptography algorithms like AES, Blowfish etc. and asymmetric key cryptography algorithms like RSA, Diffie-Hellman etc. can be used to prevent unauthorized access to the training data and ML model.
   - Homomorphic Encryption performs computation on encrypted training data [72, 73] thus ensuring the confidentiality of the training data. Works like [3] perform FHE on deep neural networks.
   - Secure Multi Party Computation is used in privacy-preserving ML [74, 75] to prevent leakage of private information from distributed ML model training.
(2) **Anonymization based** procedures mask personally identifiable information to avoid identity disclosure, thereby assuring confidentiality. Following are the main techniques available to implement anonymization.
   - Generalization involves replacing specific or sensitive data values with more generalized or abstract ranges of data values. It obscures individual identities or characteristics while retaining statistical trends and patterns in the data.
   - Suppression involves masking of specific data points or records from the dataset. This involves removing specific individuals with unique attributes or small groups that could be easily identified.
   - Decomposition involves splitting a dataset into multiple parts or tables, each containing a subset of the original data. It can enhance privacy by separating sensitive attributes from less sensitive ones, making it more difficult to link them.
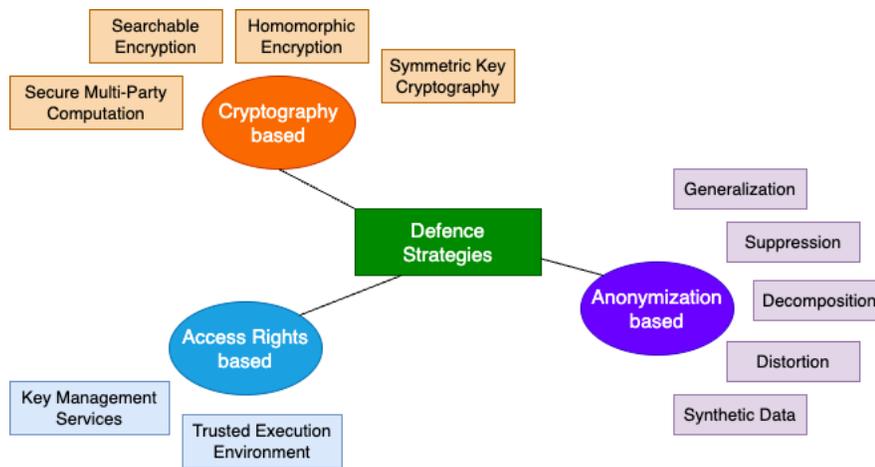


Fig. 8. Hierarchy of defences

- Interference or Distortion introduces noise or random data into the dataset to mask sensitive information. It can be performed by adding random noise or multiplying by a constant to inflate actual values. This makes it more challenging to discern individual details or patterns, even when some data remains identifiable. Differential Privacy is one way to achieve ML privacy by using such data perturbation.
- Synthetic Data can be generated via statistical techniques using median or standard deviation, or classical linear regression. Synthetic image data can also be generated using Variational Auto Encoders (VAEs), Generative Adversarial Networks (GANs) etc. These can be shared in lieu of private and sensitive images for training ML models.

(3) **Access Restrictions based** controls ensure that procedures perform authentication, authorization and maintain access control policies to regulate who has access to different parts of the system. Such enforcement can reduce the impact of risk.

- Key Management Services provide secure environments for generating cryptographic keys, and also supports Customer managed keys via 'Bring Your Own Key' initiatives to ensure consistent key management policies.
- Trusted Execution Environments create isolated, secure regions within a computing device to provide a high level of security and privacy to perform sensitive tasks.

The user is afforded the flexibility to deploy one or more of the above-mentioned defence measures, facilitating a customized implementation aligned with their distinct security requirements. Figure 8 arranges these measures hierarchically, according to the level at which the defence is required. Cryptography based defences are required at the data level, especially when the resources belonging to 1 user must be secured. Anonymization based defences apply well to groups of users, while Access rights based defences work at the application or API service level.
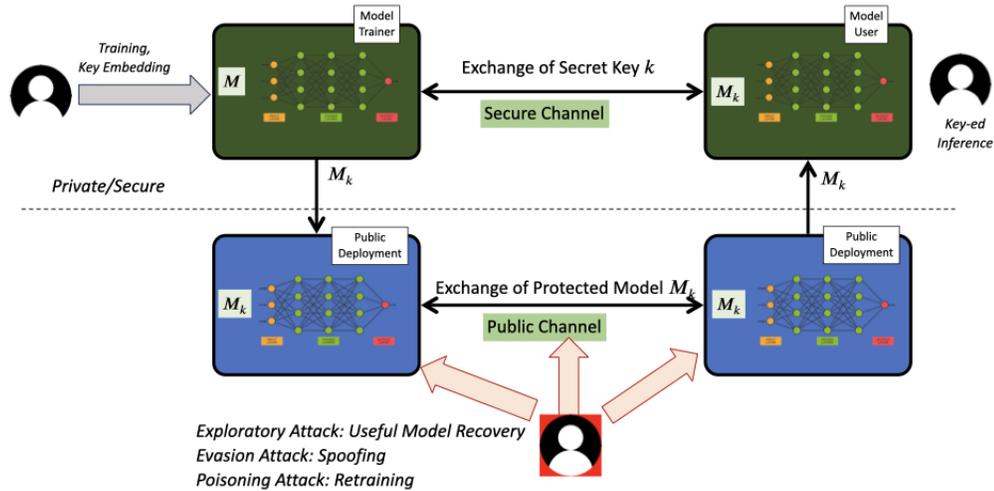


Fig. 9. Encrypting an ML model.

Table 5. Mapping known Privacy and Robustness attacks on Data and Model onto the proposed taxonomy of Section 4.

| Setting | Data Owner | Model Owner | Attacks | | | Defences |
|---|---|---|---|---|---|---|
| SDSM | 1 | 1 | Model - Privacy | Model inversion Model reconstruction Model extraction | Cryptographic<br>Access rights based | Data-in-transit encryption Compute on encrypted data Trusted Execution Environment |
| SDMM | 1 | n | Model - Privacy | Model inversion Model reconstruction Model extraction | Cryptographic<br>Access rights based | Data-in-transit encryption Compute on encrypted data Trusted Execution Environment |
| MDSM | n | 1 | Model - Privacy | Model inversion Model reconstruction Model extraction | Cryptographic<br>Access rights based | Data-in-transit encryption Compute on encrypted data Trusted Execution Environment |
| | | | Data - Privacy | Infer class representatives Infer statistical properties | Distortion | Differential Privacy |
| | | | | Infer attributes Reconstruct training input and labels Membership Inference | Cryptographic | Compute on encrypted data Symmetric Key cryptography<br><br>Data-in-transit encryption Symmetric Key cryptography |
| | | | Data - Robustness | Data poisoning | | |
| MDMM | n | m ≤ n | Model - Privacy | Model inversion Model reconstruction Model extraction | Cryptographic<br>Access rights based | Data-in-transit encryption Compute on encrypted data Trusted Execution Environment |
| | | | Data - Privacy | Infer class representatives Infer statistical properties | Distortion | Differential Privacy |
| | | | | Infer attributes Reconstruct training input and labels Membership Inference | Cryptographic | Compute on encrypted data Symmetric Key cryptography<br><br>Data-in-transit encryption Symmetric Key cryptography |
| | | | Data - Robustness | Data poisoning | | |
| | | | Model - Robustness | Model poisoning Backdoor attacks | Access rights based | Trusted Execution Environment |

Figure 9 depicts a potential defence strategy that leverages cryptographic measures to transmit the private and protected ML model over public channels. This involves secret key exchange over secure channels, designed to embed the keys into the architecture of the neural network model without disrupting training layers.

A comprehensive view of the proposed taxonomy along with recommended defence techniques is presented in Table 5 to foster a user-centric approach to ML solutioning. Consider the example of a membership inference attack on a mobile application used to predict the next word while users type on the device keyboard. Since the primary medium of information entry in mobile devices is the keyboard, it is privy to a lot of private and sensitive user information if all the keystrokes are recorded and used to train a centralized Next Word Prediction model. This environment context corresponds to the Multiple Data Single Model (MDSM) setting of the proposed taxonomy. Following Table

5, membership inference attack is a type of attack on data privacy and the recommended defence for it can be either cryptographic or distortion based, by implementing Differential Privacy.

For model consumers of AIaaS, such a table provides a user-friendly reference guide that helps them identify security gaps in their solution architecture. They can implement practical solutions or guidance to resolve these issues independently or leverage out-of-the-box services offered by the cloud provider. On the other hand, model providers can benefit from this view by creating the building blocks which consumers will need to build robust and secure applications. It grants a competitive edge to the cloud providers who have a more extensive portfolio.

## 6 CONCLUDING REMARKS

As machine learning becomes increasingly pervasive in all domains, we realize the importance of educating the general public on the involved nature of privacy and security vulnerabilities in ML applications. Recognizing the privacy risks of tools used today can also influence legislature and give more rights to the public. In a strong move towards the European Union's GDPR right to delete personal data, Amsterdam developed software that blurs people in images captured by their public street mapping cameras. Even though these images were used to train ML models, individuals' privacy is protected by concealing their face and skin colour. There is a pressing need to understand the evolving landscape of security threats and their defences in the context of AI enabled personal and critical applications.

We conducted a comprehensive study of the privacy guarantees offered by major cloud service providers and proposed a threat model to identify the systemic risks. We proposed a taxonomy to classify ML applications based on ownership of training data and usage rights of the trained model. Then we presented an extensive list of attacks possible to data and model privacy as well as robustness and proposed a classification for the defence techniques against it. Finally we unified the initial taxonomy with the different types of attacks and defences based on their characteristics.

## ACKNOWLEDGEMENT

## REFERENCES

[1] J. Yu, Y. Wu, D. Shu, M. Jin, and X. Xing, "Assessing prompt injection risks in 200+ custom gpts," 2023.

[2] A. Salem, G. Cherubin, D. Evans, B. Köpf, A. Paverd, A. Suri, S. Tople, and S. Zanella-Béguelin, "Sok: Let the privacy games begin! a unified treatment of data inference privacy in machine learning," in *2023 IEEE Symposium on Security and Privacy (SP)*, pp. 327–345, 2023.

[3] J.-W. Lee, H. Kang, Y. Lee, W. Choi, J. Eom, M. Deryabin, E. Lee, J. Lee, D. Yoo, Y.-S. Kim, and J.-S. No, "Privacy-preserving machine learning with fully homomorphic encryption for deep neural network," *IEEE Access*, vol. 10, pp. 30039–30054, 2022.

[4] H. Hu, Z. Salcic, L. Sun, G. Dobbie, P. S. Yu, and X. Zhang, "Membership inference attacks on machine learning: A survey," *ACM Comput. Surv.*, vol. 54, sep 2022.

[5] B. Liu, M. Ding, S. Shaham, W. Rahayu, F. Farokhi, and Z. Lin, "When machine learning meets privacy: A survey and outlook," *ACM Comput. Surv.*, vol. 54, mar 2021.

[6] X. Liu, L. Xie, Y. Wang, J. Zou, J. Xiong, Z. Ying, and A. V. Vasilakos, "Privacy and security issues in deep learning: A survey," *IEEE Access*, vol. 9, pp. 4566–4593, 2021.

[7] R. Shokri, M. Stronati, C. Song, and V. Shmatikov, "Membership inference attacks against machine learning models," in *2017 IEEE Symposium on Security and Privacy (SP)*, pp. 3–18, 2017.

[8] B. Hitaj, G. Ateniese, and F. Perez-Cruz, "Deep models under the gan: Information leakage from collaborative deep learning," in *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*, CCS '17, (New York, NY, USA), p. 603–618, Association for Computing Machinery, 2017.

[9] Z. Wang, M. Song, Z. Zhang, Y. Song, Q. Wang, and H. Qi, "Beyond inferring class representatives: User-level privacy leakage from federated learning," in *IEEE INFOCOM 2019 - IEEE Conference on Computer Communications*, pp. 2512–2520, 2019.

[10] C. Priebe, K. Vaswani, and M. Costa, "Enclavedb: A secure database using sgx," in *2018 IEEE Symposium on Security and Privacy (SP)*, pp. 264–278, IEEE, 2018.

[11] F. Mo, Z. Tarkhani, and H. Haddadi, "Sok: machine learning with confidential computing," *arXiv preprint arXiv:2208.10134*, 2022.

[12] A. Acar, H. Aksu, A. S. Uluagac, and M. Conti, "A survey on homomorphic encryption schemes: Theory and implementation," *ACM Comput. Surv.*, vol. 51, jul 2018.

[13] R. Curtmola, J. Garay, S. Kamara, and R. Ostrovsky, "Searchable symmetric encryption: Improved definitions and efficient constructions," in *Proceedings of the 13th ACM Conference on Computer and Communications Security*, CCS '06, (New York, NY, USA), p. 79–88, Association for Computing Machinery, 2006.

[14] P. Samarati and S. D. C. di Vimercati, "Data protection in outsourcing scenarios: Issues and directions," in *Proceedings of the 5th ACM Symposium on Information, Computer and Communications Security*, ASIACCS '10, (New York, NY, USA), p. 1–14, Association for Computing Machinery, 2010.

[15] TerryLanfear, "Protection of customer data in Azure — learn.microsoft.com." https://learn.microsoft.com/en-us/azure/security/fundamentals/protection-customer-data. [Accessed 11-09-2023].

[16] KCCross, "Technical reference details about encryption — learn.microsoft.com." https://learn.microsoft.com/en-us/purview/technical-reference-details-about-encryption. [Accessed 14-09-2023].

[17] TerryLanfear, "Double Encryption in Microsoft Azure — learn.microsoft.com." https://learn.microsoft.com/en-us/azure/security/fundamentals/double-encryption. [Accessed 14-09-2023].

[18] "Data Privacy - Amazon Web Services (AWS) — aws.amazon.com." https://aws.amazon.com/compliance/data-privacy/. [Accessed 11-09-2023].

[19] "Encrypting File Data with Amazon Elastic File System - Encrypting File Data with Amazon Elastic File System — docs.aws.amazon.com." https://docs.aws.amazon.com/whitepapers/latest/efs-encrypted-file-systems/efs-encrypted-file-systems.html. [Accessed 11-09-2023].

[20] "Google security overview | Documentation | Google Cloud — cloud.google.com." https://cloud.google.com/docs/security/overview/whitepaper. [Accessed 11-09-2023].

[21] "Huawei Cloud Security White Paper | Huawei Cloud — huaweicloud.com." https://res-static.hc-cdn.cn/cloudbu-site/intl/en-us/TrustCenter/WhitePaper/Best%20Practices/SecurityWhitepaper_intl_en.pdf, 2022. [Accessed 11-09-2023].

[22] "Huawei Cloud Data Security White Paper | Huawei Cloud — huaweicloud.com." https://res-static.hc-cdn.cn/cloudbu-site/intl/en-us/TrustCenter/WhitePaper/Best%20Practices/DataSecurityWhitepaper_intl_en.pdf, 2022. [Accessed 11-09-2023].

[23] "Encryption of data in transit — ibm.com." https://www.ibm.com/docs/en/db2/11.5?topic=encryption-data-in-transit. [Accessed 11-09-2023].

[24] "What is Confidential Computing? | IBM — ibm.com." https://www.ibm.com/topics/confidential-computing. [Accessed 11-09-2023].

[25] "Oracle Cloud Infrastructure Security Architecture | Oracle — oracle.com." https://www.oracle.com/a/ocom/docs/oracle-cloud-infrastructure-security-architecture.pdf, 2021. [Accessed 11-09-2023].

[26] "Alibaba Cloud Security White Paper | Ali Cloud — alicloud.com." https://alicloud-common.oss-ap-southeast-1.aliyuncs.com/2021/Whitepaper/Alibaba%20Cloud%20Security%20Whitepaper%20-%20International%20Edition%20V2.1%20%282021%29.pdf, 2021. [Accessed 11-09-2023].

[27] "Data protection - SageMaker Studio Administration Best Practices — docs.aws.amazon.com." https://docs.aws.amazon.com/whitepapers/latest/sagemaker-studio-admin-best-practices/data-protection.html. [Accessed 11-09-2023].

[28] "Customer-managed encryption keys (CMEK) | Vertex AI | Google Cloud — cloud.google.com." https://cloud.google.com/vertex-ai/docs/general/cmek. [Accessed 11-09-2023].

[29] S. Cheng, Y. Dong, T. Pang, H. Su, and J. Zhu, "Improving black-box adversarial attacks with a transfer-based prior," *Advances in neural information processing systems*, vol. 32, 2019.

[30] A.-M. Cretu, D. Jones, Y.-A. de Montjoye, and S. Tople, "Re-aligning shadow models can improve white-box membership inference attacks," *arXiv preprint arXiv:2306.05093*, 2023.

[31] J. Jia, A. Salem, M. Backes, Y. Zhang, and N. Z. Gong, "Memguard: Defending against black-box membership inference attacks via adversarial examples," in *Proceedings of the 2019 ACM SIGSAC conference on computer and communications security*, pp. 259–274, 2019.

[32] G. Nanfack, A. Fulleringer, J. Marty, M. Eickenberg, and E. Belilovsky, "Adversarial attacks on the interpretation of neuron activation maximization," *arXiv preprint arXiv:2306.07397*, 2023.

[33] G. Han, J. Choi, H. Lee, and J. Kim, "Reinforcement learning-based black-box model inversion attacks," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 20504–20513, 2023.

[34] I. Driouich, C. Xu, G. Neglia, F. Giroire, and E. Thomas, "Local model reconstruction attacks in federated learning and their uses," *arXiv preprint arXiv:2210.16205*, 2022.

[35] L. Lyu and C. Chen, "A novel attribute reconstruction attack in federated learning," *arXiv preprint arXiv:2108.06910*, 2021.

[36] S. Mehnaz, N. Li, and E. Bertino, "Black-box model inversion attribute inference attacks on classification models," *arXiv preprint arXiv:2012.03404*, 2020.

[37] C. Fu, X. Zhang, S. Ji, J. Chen, J. Wu, S. Guo, J. Zhou, A. X. Liu, and T. Wang, "Label inference attacks against vertical federated learning," in *31st USENIX Security Symposium (USENIX Security 22)*, pp. 1397–1414, 2022.

[38] C. A. Choquette-Choo, F. Tramer, N. Carlini, and N. Papernot, "Label-only membership inference attacks," in *International conference on machine learning*, pp. 1964–1974, PMLR, 2021.

[39] Z. Wang, M. Song, Z. Zhang, Y. Song, Q. Wang, and H. Qi, "Beyond inferring class representatives: User-level privacy leakage from federated learning," in *IEEE INFOCOM 2019-IEEE conference on computer communications*, pp. 2512–2520, IEEE, 2019.

[40] J. Lin, L. Dang, M. Rahouti, and K. Xiong, "Ml attack models: adversarial attacks and data poisoning attacks," *arXiv preprint arXiv:2112.02797*, 2021.

[41] Z. Lin, Y. Shi, and Z. Xue, "Idsgan: Generative adversarial networks for attack generation against intrusion detection," in *Pacific-asia conference on knowledge discovery and data mining*, pp. 79–91, Springer, 2022.

[42] G. Liu and L. Lai, "Provably efficient black-box action poisoning attacks against reinforcement learning," *Advances in Neural Information Processing Systems*, vol. 34, pp. 12400–12410, 2021.

[43] C. Zhu, W. R. Huang, H. Li, G. Taylor, C. Studer, and T. Goldstein, "Transferable clean-label poisoning attacks on deep neural nets," in *International Conference on Machine Learning*, pp. 7614–7623, PMLR, 2019.

[44] A. R. Shahid, A. Imteaj, P. Y. Wu, D. A. Igoche, and T. Alam, "Label flipping data poisoning attack against wearable human activity recognition system," in *2022 IEEE Symposium Series on Computational Intelligence (SSCI)*, pp. 908–914, IEEE, 2022.

[45] T. B. Brown and D. Mané, "Aurko roy, martín abadi, and justin gilmer," *Adversarial patch*, vol. 2, p. 1, 2017.

[46] V. Shejwalkar, A. Houmansadr, P. Kairouz, and D. Ramage, "Back to the drawing board: A critical evaluation of poisoning attacks on production federated learning," in *2022 IEEE Symposium on Security and Privacy (SP)*, pp. 1354–1371, IEEE, 2022.

[47] J. Dumford and W. Scheirer, "Backdooring convolutional neural networks via targeted weight perturbations," in *2020 IEEE International Joint Conference on Biometrics (IJCB)*, pp. 1–9, IEEE, 2020.

[48] C. Szegedy, W. Zaremba, I. Sutskever, J. Bruna, D. Erhan, I. Goodfellow, and R. Fergus, "Intriguing properties of neural networks," *arXiv preprint arXiv:1312.6199*, 2013.

[49] B. Biggio and F. Roli, "Wild patterns: Ten years after the rise of adversarial machine learning," in *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security*, CCS '18, (New York, NY, USA), p. 2154–2156, Association for Computing Machinery, 2018.

[50] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. u. Kaiser, and I. Polosukhin, "Attention is all you need," in *Advances in Neural Information Processing Systems* (I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, eds.), vol. 30, Curran Associates, Inc., 2017.

[51] F. Zhuang, Z. Qi, K. Duan, D. Xi, Y. Zhu, H. Zhu, H. Xiong, and Q. He, "A comprehensive survey on transfer learning," *Proceedings of the IEEE*, vol. 109, no. 1, pp. 43–76, 2021.

[52] T. Li, A. K. Sahu, A. Talwalkar, and V. Smith, "Federated learning: Challenges, methods, and future directions," *IEEE Signal Processing Magazine*, vol. 37, no. 3, pp. 50–60, 2020.

[53] T. Khan, K. Nguyen, and A. Michalas, "Split ways: Privacy-preserving training of encrypted data using split learning," 2023.

[54] M. Fredrikson, E. Lantz, S. Jha, S. Lin, D. Page, and T. Ristenpart, "Privacy in pharmacogenetics: An End-to-End case study of personalized warfarin dosing," in *23rd USENIX Security Symposium (USENIX Security 14)*, (San Diego, CA), pp. 17–32, USENIX Association, Aug. 2014.

[55] P. W. Koh, J. Steinhardt, and P. Liang, "Stronger data poisoning attacks break data sanitization defenses," *Machine Learning*, pp. 1–47, 2022.

[56] D. Puccinelli and M. Haenggi, "Wireless sensor networks: applications and challenges of ubiquitous sensing," *IEEE Circuits and Systems Magazine*, vol. 5, no. 3, pp. 19–31, 2005.

[57] P. Asghari, A. M. Rahmani, and H. H. S. Javadi, "Internet of things applications: A systematic review," *Computer Networks*, vol. 148, pp. 241–261, 2019.

[58] L. Nie, X. Wang, S. Wang, Z. Ning, M. S. Obaidat, B. Sadoun, and S. Li, "Network traffic prediction in industrial internet of things backbone networks: A multitask learning mechanism," *IEEE Transactions on Industrial Informatics*, vol. 17, no. 10, pp. 7123–7132, 2021.

[59] T. Yu, T. Li, Y. Sun, S. Nanda, V. Smith, V. Sekar, and S. Seshan, "Learning context-aware policies from multiple smart homes via federated multi-task learning," in *2020 IEEE/ACM Fifth International Conference on Internet-of-Things Design and Implementation (IoTDI)*, pp. 104–115, 2020.

[60] F. Chang, J. Dean, S. Ghemawat, W. C. Hsieh, D. A. Wallach, M. Burrows, T. Chandra, A. Fikes, and R. E. Gruber, "Bigtable: A distributed storage system for structured data," *ACM Trans. Comput. Syst.*, vol. 26, jun 2008.

[61] C. J. Hoofnagle, B. van der Sloot, and F. Z. Borgesius, "The european union general data protection regulation: what it is and what it means," *Information & Communications Technology Law*, vol. 28, no. 1, pp. 65–98, 2019.

[62] W. B. Chik, "The singapore personal data protection act and an assessment of future trends in data privacy reform," *Computer Law & Security Review*, vol. 29, no. 5, pp. 554–575, 2013.

[63] M. C. Addis and M. Kutar, "The general data protection regulation (gdpr), emerging technologies and uk organisations: Awareness, implementation and readiness," *UK Academy for Information Systems Conference Proceedings 2018. 29*, 2018.

[64] B. Gaff, H. Sussman, and J. Geetter, "Privacy and big data," *Computer*, vol. 47, pp. 7–9, 06 2014.

[65] Z. Zhang, P. Luo, C. C. Loy, and X. Tang, "Facial landmark detection by deep multi-task learning," in *Computer Vision – ECCV 2014* (D. Fleet, T. Pajdla, B. Schiele, and T. Tuytelaars, eds.), (Cham), pp. 94–108, Springer International Publishing, 2014.

[66] S. Ö. Arık, M. Chrzanowski, A. Coates, G. Diamos, A. Gibiansky, Y. Kang, X. Li, J. Miller, A. Ng, J. Raiman, S. Sengupta, and M. Shoeybi, "Deep voice: Real-time neural text-to-speech," in *Proceedings of the 34th International Conference on Machine Learning* (D. Precup and Y. W. Teh, eds.), vol. 70 of *Proceedings of Machine Learning Research*, pp. 195–204, PMLR, 06–11 Aug 2017.

[67] M. Maqsood, F. Nazir, U. Khan, F. Aadil, H. Jamal, I. Mehmood, and O.-y. Song, "Transfer learning assisted classification and detection of alzheimer's disease stages using 3d mri scans," *Sensors*, vol. 19, no. 11, 2019.

[68] A. G. Roy, S. Siddiqui, S. Pölsterl, N. Navab, and C. Wachinger, "Braintorrent: A peer-to-peer environment for decentralized federated learning," 2019.

[69] C. T. Dinh, N. Tran, and J. Nguyen, "Personalized federated learning with moreau envelopes," in *Advances in Neural Information Processing Systems* (H. Larochelle, M. Ranzato, R. Hadsell, M. Balcan, and H. Lin, eds.), vol. 33, pp. 21394–21405, Curran Associates, Inc., 2020.

[70] Y. Jiang, J. Konečný, K. Rush, and S. Kannan, "Improving federated learning personalization via model agnostic meta learning," *CoRR*, vol. abs/1909.12488, 2019.

[71] V. Smith, C.-K. Chiang, M. Sanjabi, and A. S. Talwalkar, "Federated multi-task learning," in *Advances in Neural Information Processing Systems* (I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, eds.), vol. 30, Curran Associates, Inc., 2017.

[72] B. Reagen, W.-S. Choi, Y. Ko, V. T. Lee, H.-H. S. Lee, G.-Y. Wei, and D. Brooks, "Cheetah: Optimizing and accelerating homomorphic encryption for private inference," in *2021 IEEE International Symposium on High-Performance Computer Architecture (HPCA)*, pp. 26–39, 2021.

[73] N. Chakraborty and G. Patra, "Functional encryption for secured big data analytics," *International Journal of Computer Applications*, vol. 107, no. 16, 2014.

[74] C. zhi Gao, Q. Cheng, P. He, W. Susilo, and J. Li, "Privacy-preserving naive bayes classifiers secure against the substitution-then-comparison attack," *Information Sciences*, vol. 444, pp. 72–88, 2018.

[75] X. S. Wang, Y. Huang, Y. Zhao, H. Tang, X. Wang, and D. Bu, "Efficient genome-wide, privacy-preserving similar patient query based on private edit distance," in *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security*, CCS '15, (New York, NY, USA), p. 492–503, Association for Computing Machinery, 2015.