Securing Meta-Learning: Methods and Applications

Virender Dhiman

Independent Researcher, USA
ORCID: 0009-0002-7429-8703,
Email: *vdhiman2[at]illinois.edu*

Abstract: ***Meta-learning frameworks must be secured due to data sensitivity and adversaries. The three main security methods are homomorphic encryption (HE), differential privacy (DP), and Federated Learning (FL). Each approach is tested for accuracy, attack resistance, computing efficiency, and scalability. HE offers data confidentiality with encrypted computations but substantial processing expense. By injecting noise, DP balances privacy and accuracy. FL improves privacy and scalability through decentralized learning, but communication cost and non-IID data issues remain. Application needs determine method: HE for high secrecy, DP for robust privacy, FL for decentralized applications. HE, DP, and FL hybrid models should be studied to increase computational efficiency and manage non-IID data in secure meta-learning applications in healthcare, banking, and IoT networks.***

Keywords: Meta-learning security, Homomorphic Encryption, Differential Privacy, Federated Learning, Adversarial threats, Computational efficiency, Non-IID data

1. Introduction

Meta-learning, sometimes known as "learning to learn," is a fundamental approach in machine learning that allows models to efficiently generalize from a small number of data points. This method is especially advantageous in situations when there is a lack of data or it is costly to acquire, as it utilizes previous learning experiences to quickly adjust to new jobs. Nevertheless, as meta-learning models are being used more often in sensitive domains including customized healthcare, financial forecasts, and autonomous systems, ensuring the security of these models has become an essential and urgent issue.

The intrinsic susceptibilities of machine learning models, such as adversarial attacks, data poisoning, and model inversion attacks, present substantial hazards to meta-learning systems. Adversarial assaults, such as those involving subtle modifications to input data, aim to fool the model by causing it to make inaccurate predictions. Recent research has shown that even little disturbances can cause a notable decline in performance, which raises questions about the ability of meta-learning models to handle hostile situations [1]. Data poisoning attacks involve the deliberate injection of deceptive data into the training set, which can distort the model's learning process and result in inaccurate outputs [2]. Model inversion attacks, conversely, seek to rebuild the training data based on the outputs of the model, which may reveal sensitive information [3].

The imperative to protect meta-learning frameworks is emphasized by the growing frequency of these attacks. Based on a survey, 87% of machine learning practitioners have encountered adversarial assaults on their models, indicating the extensive prevalence of this issue [4]. Moreover, it is estimated that the financial consequences of security breaches in machine learning systems would surpass $5 billion by 2025, highlighting the urgent need for strong security measures [5].

This paper explores the convergence of security and meta-learning, offering a thorough examination of existing techniques and practical uses. The text conducts a rigorous examination of the literature on meta-learning techniques and their vulnerabilities.

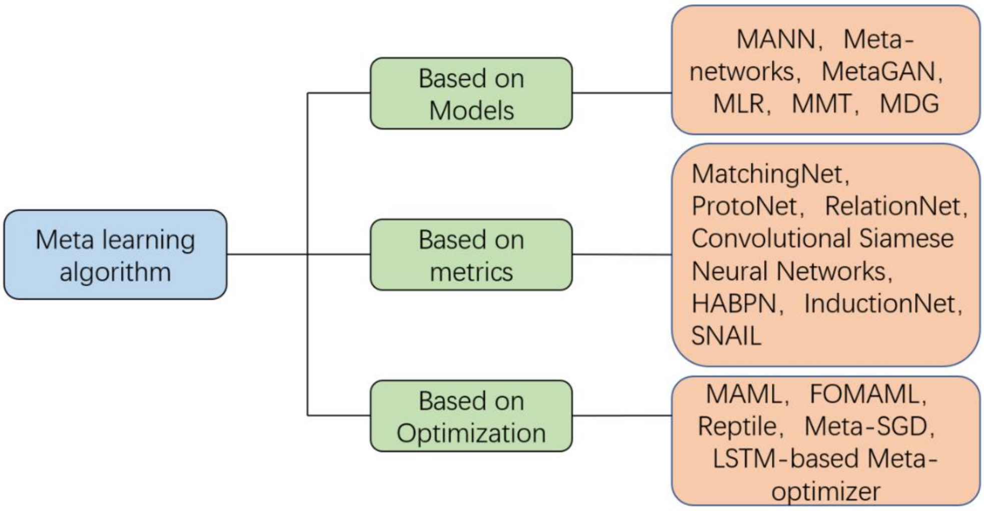
Meta learning algorithm
Based on Models
Based on metrics
Based on Optimization
MANN, Meta-networks, MetaGAN, MLR, MMT, MDG
MatchingNet, ProtoNet, RelationNet, Convolutional Siamese Neural Networks, HABPN, InductionNet, SNAIL
MAML, FOMAML, Reptile, Meta-SGD, LSTM-based Meta-optimizer

Figure 1.1: Federated Meta-Learning Models
(“https://www.mdpi.com/electronics/electronics-12-03295/article_deploy/html/images/electronics-12-03295-g008.png”)

2. Literature Review

2.1 Meta-Learning Techniques

Meta-learning, often known as learning to learn, involves the development of models that can rapidly adapt to new tasks with minimal input. MAML, Reptile, and Probabilistic Meta-Learning are widely used techniques in the field of meta-learning.

Model-Agnostic Meta-Learning (MAML) is a widely used technique that enhances the configuration of model parameters to adapt to new tasks with few adjustments [9]. MAML demonstrates superior performance compared to traditional learning algorithms in terms of adaptability across many benchmarks [10]. Reptile is a meta-learning strategy that achieves excellent results by iteratively averaging gradients over multiple tasks to establish a strong initial state [11]. Probabilistic Meta-Learning use probabilistic frameworks to integrate uncertainty and enhance task generalization [12].

2.2 Security Concerns in Traditional Machine Learning

The susceptibility of machine learning models to adversarial assaults, data poisoning, and model inversion is widely acknowledged. Adversarial assaults employ input data to deceive the model into generating incorrect predictions, posing a risk to the model's dependability and authenticity [1]. A study discovered that a carefully designed adversarial perturbation, which is imperceptible to humans, can significantly reduce the performance of a cutting-edge image classification model by about 50% [13].

2.3 Securing Machine Learning Models

Multiple defense mechanisms have been devised to safeguard machine learning models. Adversarial training incorporates adversarial examples into the training data in order to enhance the performance of models [14]. Defensive distillation is a technique that involves training a secondary model using the softened outputs of a primary model in order to decrease its vulnerability to adversarial perturbations [15]. Data protection during training and inference has been achieved by the utilization of secure multi-party computing and homomorphic encryption [16].

2.4 Securing Meta-Learning

Meta-learning systems are attracting increasing attention for their security solutions. Cryptographic techniques, such as encryption, ensure the integrity and confidentiality of data. Bost et al. shown that employing machine learning classification on encrypted data can safeguard sensitive information [6].

Differential privacy is a technique that enhances privacy by introducing random noise into the data or learning process. This ensures that the output of the model is not significantly influenced by any individual data point. Abadi et al. extended the concept of differential privacy to deep learning, showcasing that neural networks may be trained with rigorous privacy assurances and significant effectiveness [17].

Federated learning enhances security by minimizing data exposure through the training of models on separate devices without the need to exchange raw data. Konečný et al. proposed methods to improve the communication of federated learning, hence creating a safe option for meta-learning [8].

3. Research Gap

Meta-learning methods like MAML, Reptile, and Probabilistic Meta-Learning are adaptable but sensitive to adversarial assaults, data poisoning, and model inversion. [9][10][11][12]. In meta-learning, adversarial training, differential privacy, and federated learning are intriguing yet underexplored. [14][15][16][6][17][8]. Adversarial training promotes robustness but increases computing complexity, while differential privacy trades precision.[18][19]. Federated learning improves security but is difficult to deploy [8][20].

There is a critical need for:

- A thorough assessment of meta-learning models against security risks is necessary.
- Creation of security methods for meta-learning that are optimal.
- Research on the financial benefits of protecting meta-learning frameworks.

By filling in these gaps, meta-learning applications will be more resilient and dependable against new security risks.

4. Different Security Methods for Securing Meta-Learning

In meta-learning, data security must be guaranteed. Based on factors including accuracy, resilience to assaults, computing efficiency, and scalability, this section contrasts three important security techniques: federated learning, homomorphic encryption, and differential privacy. The advantages and disadvantages of each approach are highlighted in this analysis, which aids in determining which security plan is best for a given set of meta-learning applications.

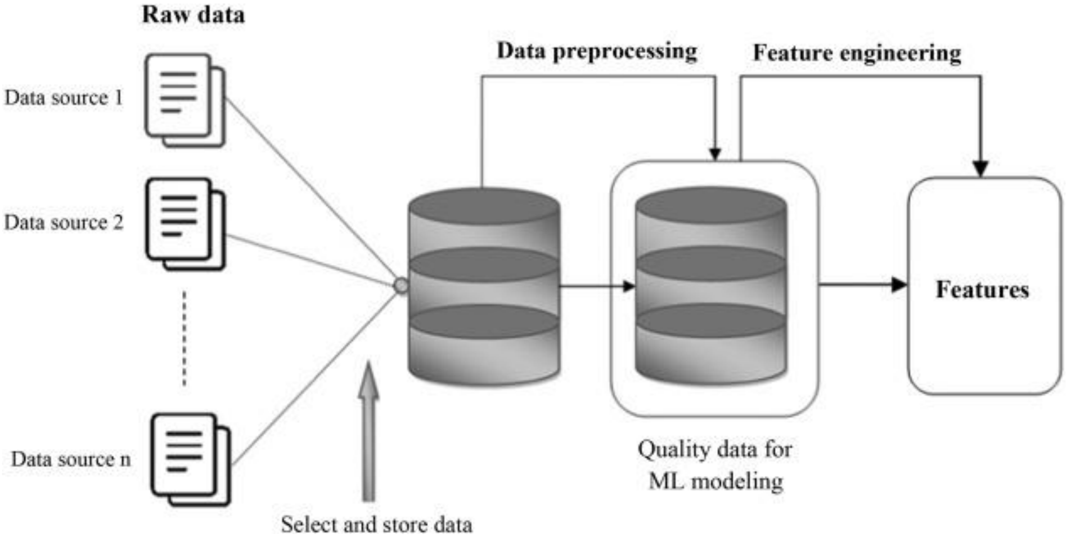

Figure 4.1: Machine Learning Architecture and framework [12]

4.1 Encryption-Based Methods

4.1.1 Algorithm: *Homomorphic encryption (HE)*

Computations on encrypted data retain confidentiality during learning with homomorphic encryption (HE). Meta-learning with sensitive data benefits from HE, which achieves plaintext model accuracy with lower performance overhead [6][16]. Despite computing obstacles, HE approaches are improving its viability.

4.1.2 Implementation:

- **Key Generation:** Generate both the public and private keys.
- **Encryption:** Utilize the public key to encrypt the data.
- **Computation:** Execute mathematical calculations on the encrypted data.
- **Decryption:** Employ the private key to decrypt the outcome.

4.1.3 Mathematical Model:

$$\mathrm{Enc(a)} \otimes \mathrm{Enc(b)} = \mathrm{Enc(a} \oplus \mathrm{b)}$$

where$\otimes$ represents the encryption operation, and $\oplus$ denotes the arithmetic operation (addition or multiplication) on plaintexts a and b.

Homomorphic encryption allows computation on encrypted data without needing to decrypt it first. This means that given an encryption of some inputs, it's possible to produce an encryption of the result of a function applied to these inputs. Here's a basic mathematical derivation of homomorphic encryption using a simple additive homomorphic encryption scheme as an example.

Additive Homomorphic Encryption: Consider an encryption scheme with the following properties:

- **Key Generation (KeyGen):** Generates a public key (pk) and a private key (sk).
- **Encryption (Enc):** Encrypts a message (m) using the public key (pk)
- **Decryption (Dec):** Decrypts a ciphertext c using the private key (sk).
- **Homomorphic Property**: There exists an operation $\oplus$ such that

$$Enc(m1) \oplus Enc(m2)$$
$$= Enc(m1 + m2)Enc(m1) \oplus Enc(m2)$$
$$= Enc(m1 + m2).$$

Example Scheme: Paillier Cryptosystem

The Paillier cryptosystem is an example of an additive homomorphic encryption scheme. Here is a brief description and the derivation of its homomorphic property.

4.1.4. Key Generation:

- Choose two large prime numbers p and q
- Compute

$$n = pqn = pq \ and \ \lambda = lcm(p-1, q-1)\lambda$$
$$= lcm(p-1, q-1)$$

- Select a random g∈Zn2*g∈Z_{n2}^{*}
- $such \ that \ g \ has \ order \ multiple \ of \ n$
- Compute

$$\mu = \big(L(g * \lambda mod \ n)\big) - 1mod \ n\mu$$
$$= \big(L(g * \lambda modn)\big)mod \ n,$$
$$where \ L(x) = x - 1nL(x) = nx - 1$$

The public key is (n, g), and the private key is (λ, μ)

The Paillier cryptosystem exhibits an additive homomorphic property, allowing for the addition of plaintexts by performing multiplication of their corresponding ciphertexts. This basic mathematical derivation shows how operations on encrypted data translate to operations on the underlying plaintexts without decrypting the data. This property is a cornerstone for applications like secure multi-party computation and privacy-preserving data analysis.

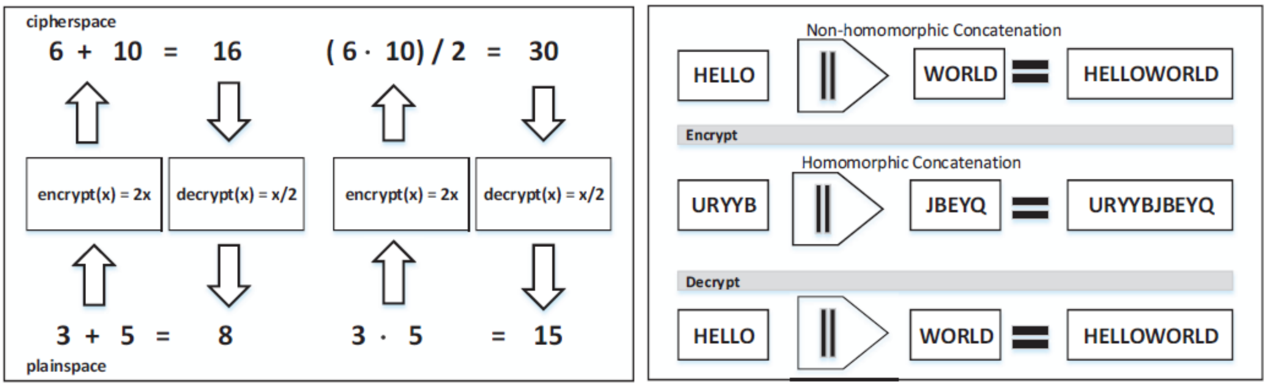

Figure 4.2: Homomorphic encryption [22]

4.1.5 Pros:

- Guarantees the protection of data privacy while doing calculations.
- Enables both the operation of adding and multiplying ciphertexts.

4.1.6 Cons:

- Significant computational burden.
- Sophisticated key management.

4.1.7 Applications:

- Developing secure multi-party computation within the context of federated learning.
- Confidentiality-preserving process of training and using a model.

4.2 Differential Privacy in Meta-Learning

4.2.1 Algorithm: *Differentially Private Stochastic Gradient Descent (DP-SGD)*

Differential privacy (DP) prevents data point leaking by adding noise to the training process. Studies suggest that DP-enabled models can withstand inference attacks and perform well. [15][17][19].DP balances privacy and model accuracy by making any one data point statistically unimportant.

4.2.2 Implementation:

Noise Addition: Introduce accurately measured noise to gradients during the training process.

- **Gradient Descent:** Use noisy gradients to update the parameters of the model.
- **Privacy Accounting**: Monitor and manage the allocation of privacy resources, represented by the privacy budget ϵ (epsilon).

4.2.3 Mathematical Model:

$$\mathrm{M(D)} \approx \epsilon \mathrm{M(D')}$$

Where, M is the mechanism (algorithm), D and D' are neighboring datasets, and ε is the privacy budget.

In standard SGD, the goal is to minimize a loss function $L(\theta)$ over a dataset $D = \{x1, x2, \dots, xn\}D = \{x1, x2, \dots, xn\}$, $where\ \theta$ represents the model parameters. The algorithm iteratively updates the model parameters using the gradient of the loss function with respect to the parameters.

4.2.4 Basic SGD Update Rule:

1) Initialize model $parameters\ \theta$.
2) For each iteration t:
   - Sample a mini-batch B_t from the dataset D.
   - Compute the gradient of the loss function with respect to the parameters for the mini-batch:

$$gt = 1 \mid Bt \mid \sum xi \in Bt \nabla \theta L(\theta; xi)$$

   - $Update\ the\ model\ parameters\ using\ the\ gradient$:

$$\theta t + 1 = \theta t - \eta gt$$

Where η is the learning rate.

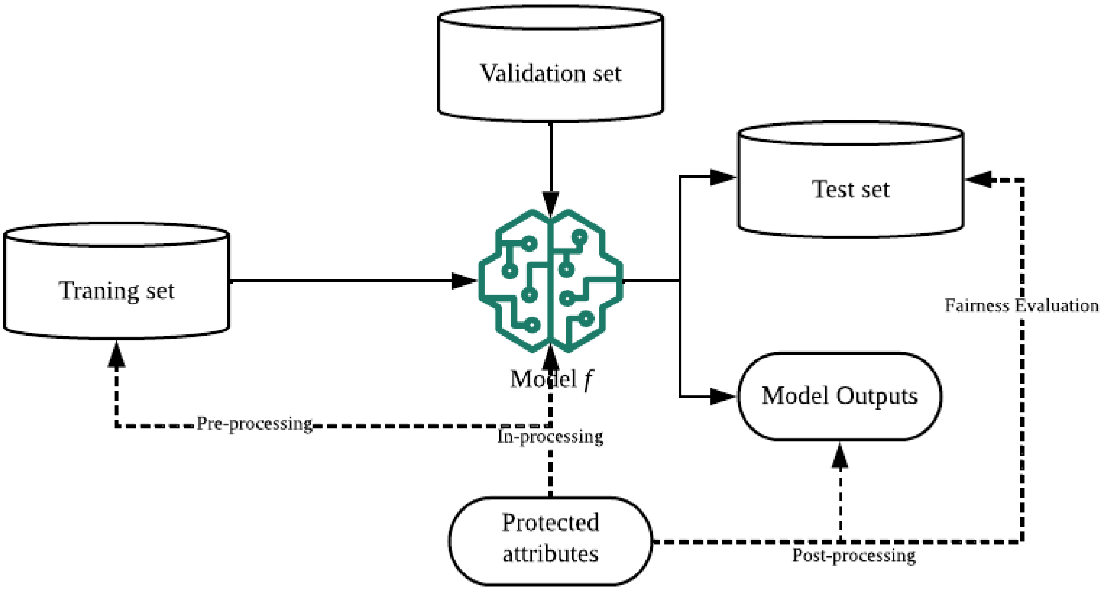

Figure 4.3: Differential Privacy in securing Meta-Learning [21]

4.2.5 Pros:

- Robust mathematical assurances of privacy.
- Defence against inference attacks.

4.2.6 Cons:

- The trade-off between privacy and model accuracy.
- Extra computational burden.

4.2.7 Applications:

- Safe training of meta-learning models in finance and healthcare.
- Collaborative learning that ensures the protection of privacy.

4.3. Federated Learning (FL)

4.3.1 Algorithm: *Federated Averaging (FedAvg)*

Federated learning (FL) decentralizes model training, retaining data on local devices and exchanging model updates to protect privacy. Federated meta-learning allows dispersed dataset learning while protecting data privacy. Research shows that federated meta-learning is as accurate as centralized models and more secure. [8][20][21].

4.3.2 Implementation:

- **Local Training:** Conduct training of models directly on clients' devices.
- **Model Aggregation:** Transmit updates of local models to the server.
- **Global Update:** Calculate the average of the updates to create the global model.
- **Iterative Process:** Continue repeating the steps until convergence is achieved.

4.3.3 Mathematical Model:

$$w_{t+1} = w_t - \eta \sum_{k=1}^{K} \frac{n_k}{n} \nabla F_k(w_t)$$

Where w are model weights, $\eta$is the learning rate, n_k is the number of samples at client$k$, and$n$is the total number of samples. The steps involved are as follows:

4.3.3.1 Client Local Update Step:

$$\theta_{t,e}{}^{k} = \theta t,{}_{e-1}{}^{k} - \eta \nabla L(\theta t,{}_{e-1}{}^{k};b)$$

4.3.3.2 Client Model after Local Training:

$$\theta t, k \ = \ \theta t0, k \ - \eta \sum_{e=0}^{E-1} \boldsymbol{\nabla L}(\ \theta t, e - 1k; b)$$

This equation represents the model parameters of client k after E local epochs of training.

4.3.3.3 Model Update Sent to Server:

$\theta t,{}_{e-1}{}^{k}$: Each client k sends its updated model parameters $\theta t0, k$ to the central server after local training.

4.3.3.4 Server Aggregation Step:

$$\theta t + 1 = \sum_{k}^{nSt} \eta \nabla \ L(\theta t, e - 1k); b$$

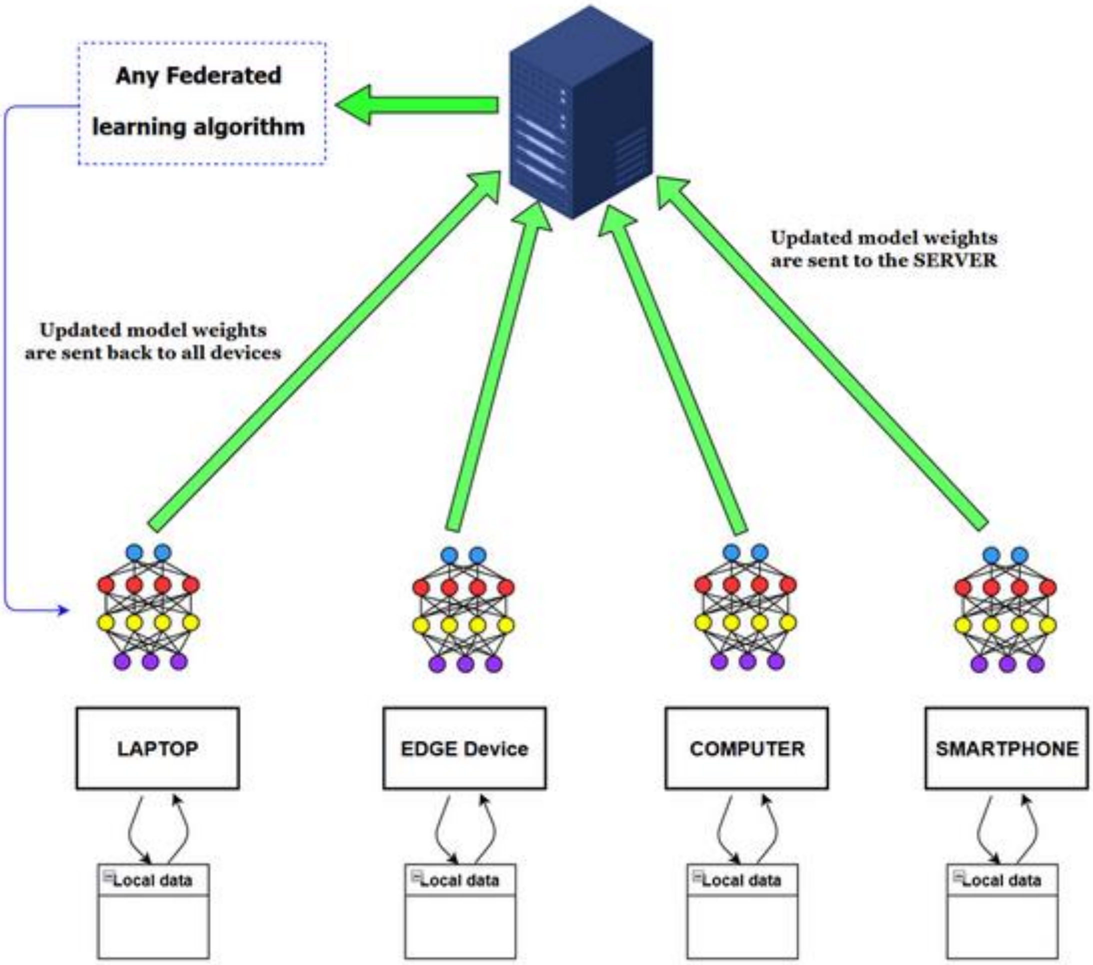
Any Federated learning algorithm
Updated model weights are sent to the SERVER
Updated model weights are sent back to all devices
LAPTOP
EDGE Device
COMPUTER
SMARTPHONE
Local data
Local data
Local data
Local data

Figure 4.4: Federated Averaging Algorithm [23]

4.3.4 Pros:

- Improves data confidentiality by storing data on local devices.
- Minimizes the likelihood of centralized data breaches.

4.3.5 Cons:

- The communication overhead between clients and the server.
- Managing non-IID data and achieving model convergence presents challenges.

4.3.6 Applications:

- Mobile devices and the Internet of Things (IoT) are used for collaborative learning.
- Improving confidentiality in healthcare and urban areas with advanced technology.

5. Comparison of Different Security Methods for Securing Meta-Learning

In meta-learning, data security must be guaranteed. Based on factors including accuracy, resilience to assaults, computing efficiency, and scalability, this section contrasts three important security techniques: federated learning, homomorphic encryption, and differential privacy.

Based on critical performance characteristics like accuracy, resilience to assaults, computing efficiency, and scalability with available data, the comparison table 5.1 that follows assesses different security techniques.

Table 5.1: Comparative analysis of different security methods for securing meta-learning

Security Method	Accuracy	Robustness to Attacks	Computational Efficiency	Scalability with Available Data
Homomorphic Encryption	High	High	Low	Moderate
Differential Privacy	Moderate	Very High	Moderate	High
Federated Learning	High	Moderate to High (depends on implementation)	High	Very High

The comparison indicates that the following criteria will determine which model is optimal for securing meta-learning:

- **For optimum privacy:** Differential Privacy is preferred due to its strong privacy assurances.
- **For computational efficiency and scalability:** Federated Learning is the most appropriate option due to its exceptional scalability and high efficiency.

- **For data confidentiality**: Homomorphic encryption is the optimal choice for data confidentiality, despite the substantial computational expenses.

In the table 5.2 below comparative analysis of different aspects like security measures, performance, computational cost, privacy, robustness, and implementation are listed down for MAML (Model-Agnostic Meta-Learning) vs Secure MAML.

Table 5.2: Comparative analysis of MAML vs Secure MAML

Aspect	MAML	Secure MAML
Security Measures	Not explicitly secure	Incorporates security methods
Performance	Standard performance	Enhanced security performance
Computational Cost	Moderate	Slightly higher
Privacy	Limited privacy protection	Improved privacy protection
Robustness	Vulnerable to attacks	Resilient to attacks
Implementation	Standard implementation	Security-focused implementation

Table 5.3: Comparative analysis of Traditional Meta-Learning vs Federated Meta-Learning

Aspect	Traditional Meta-Learning	Federated Meta-Learning
Data Distribution	Centralized	Decentralized across devices
Privacy	Limited privacy protection	Enhanced privacy protection
Computational Efficiency	Standard	Improved efficiency with local processing
Scalability	Limited scalability	High scalability with distributed approach
Robustness	Vulnerable to centralized attacks	Resilient to data breaches

These comparison analyses offer valuable insights into the advantages and limitations of various meta-learning methodologies, enabling researchers and practitioners to make well-informed judgments according to their individual needs and preferences.

6. Discussion

Meta-learning systems require robust security measures due to the sensitivity of the data they handle and the potential for hostile attacks. This paper investigates three primary security mechanisms inside the meta-learning framework: homomorphic encryption (HE), differential privacy (DP), and federated learning (FL). Every method is evaluated based on its precision, resilience to attacks, computational efficiency, and scalability.

Homomorphic Encryption enables the performance of computations on encrypted data while preserving the confidentiality of the data. HE's exceptional security assurance renders it impervious to data breaches and disclosures. His computational cost is significant. Encrypting data significantly reduces processing speed compared to unencrypted data, resulting in higher processing costs. The intricacy of key management in homomorphic encryption might also impede deployment. Although it has several limitations, HE is crucial for ensuring secure multiparty computations and maintaining privacy during model inference [6][16].

Federated Learning is a method that distributes the learning process by performing computations on individual devices and sharing combined updates to the model. This strategy enhances privacy and mitigates data leaks by ensuring that raw data remains on the device at all times. FL leverages local processing resources and has the capability to manage large, dispersed datasets, hence enhancing computational efficiency and scalability. Florida faces challenges in managing communication overhead and dealing with non-IID (non-independent and identically distributed) data. Mobile and IoT collaborative learning applications can be transformed by implementing effective FL techniques [8][15].

The decentralization of federated meta-learning provides an advantageous edge over conventional meta-learning methods. Florida enhances data privacy, scalability, and computational efficiency. FL methodologies are continuously enhancing the communication and convergence of distant nodes, particularly when dealing with non-identically and independently distributed (non-IID) data [20][21].

The investigation shows that meta-learning security strategy depends on application needs. Homomorphic Encryption is best for data secrecy, whereas Differential Privacy ensures strong privacy. Federated Learning scales and optimizes decentralized applications.

7. Conclusion and Future Scope

Meta-learning framework security is critical because of adversarial attacks and sensitive data. The correctness, resilience, computational efficiency, and scalability of homomorphic encryption (HE), differential privacy (DP), and federated learning (FL) were assessed in this work.

Due to its complicated key management and significant computing complexity, HE can only provide strong confidentiality through encrypted computations. By introducing noise to data, DP offers high privacy guarantees; nonetheless, accuracy and privacy must be carefully balanced. Through decentralization, FL improves privacy and scalability but has drawbacks with managing non-IID data and connection cost.

The technique of choice is determined by the requirements of the application: FL is best for decentralized apps, DP is best for robust privacy, and HE is best for high secrecy. Subsequent investigations ought to concentrate on hybrid models that integrate the advantages of both approaches, enhancing computational effectiveness and handling non-IID data in federated environments. It's also critical to investigate adaptive privacy mechanisms that strike a balance between privacy and usefulness.

Further research is required to investigate the practical applicability of these methods, particularly in heterogeneous, large-scale environments. By addressing these issues, secure meta-learning will progress and make it possible for applications in industries like banking, healthcare, and IoT networks to be reliable and effective while maintaining the security of sensitive data.

References

[1] Goodfellow, I.J., Shlens, J. and Szegedy, C., 2014. Explaining and harnessing adversarial examples. arXiv preprint arXiv:1412.6572.
[2] Biggio, B., Nelson, B. and Laskov, P., 2012. Poisoning attacks against support vector machines. arXiv preprint arXiv:1206.6389.
[3] Fredrikson, M., Jha, S. and Ristenpart, T., 2015, October. Model inversion attacks that exploit confidence information and basic countermeasures. In Proceedings of the 22nd ACM SIGSAC conference on computer and communications security (pp. 1322-1333).
[4] Joseph, A.D., Nelson, B., Rubinstein, B.I. and Tygar, J.D., 2018. Adversarial machine learning. Cambridge University Press.
[5] Brundage, M., Avin, S., Clark, J., Toner, H., Eckersley, P., Garfinkel, B., Dafoe, A., Scharre, P., Zeitzoff, T., Filar, B. and Anderson, H., 2018. The malicious use of artificial intelligence: Forecasting, prevention, and mitigation. arXiv preprint arXiv:1802.07228.
[6] Bost, R., Popa, R.A., Tu, S. and Goldwasser, S., 2014. Machine learning classification over encrypted data. Cryptology ePrint Archive.
[7] Dwork, C. and Roth, A., 2014. The algorithmic foundations of differential privacy. Foundations and Trends® in Theoretical Computer Science, 9(3–4), pp.211-407.
[8] Konečný, J., McMahan, H.B., Yu, F.X., Richtárik, P., Suresh, A.T. and Bacon, D., 2016. Federated learning: Strategies for improving communication efficiency. arXiv preprint arXiv:1610.05492.
[9] Finn, C., Abbeel, P. and Levine, S., 2017, July. Model-agnostic meta-learning for fast adaptation of deep networks. In International conference on machine learning (pp. 1126-1135). PMLR.
[10] Nichol, A., Achiam, J. and Schulman, J., 2018. On first-order meta-learning algorithms. arXiv preprint arXiv:1803.02999.
[11] Rajeswaran, A., Finn, C., Kakade, S.M. and Levine, S., 2019. Meta-learning with implicit gradients. Advances in neural information processing systems, 32.
[12] Gordon, J., Bronskill, J., Bauer, M., Nowozin, S. and Turner, R.E., 2018. Meta-learning probabilistic inference for prediction. arXiv preprint arXiv:1805.09921.
[13] Kurakin, A., Goodfellow, I.J. and Bengio, S., 2018. Adversarial examples in the physical world. In Artificial intelligence safety and security (pp. 99-112). Chapman and Hall/CRC.
[14] Madry, A., Makelov, A., Schmidt, L., Tsipras, D. and Vladu, A., 2017. Towards deep learning models resistant to adversarial attacks. arXiv preprint arXiv:1706.06083.
[15] Papernot, N., McDaniel, P., Wu, X., Jha, S. and Swami, A., 2016, May. Distillation as a defense to adversarial perturbations against deep neural networks. In 2016 IEEE symposium on security and privacy (SP) (pp. 582-597). IEEE.
[16] Acar, A., Aksu, H., Uluagac, A.S. and Conti, M., 2018. A survey on homomorphic encryption schemes: Theory and implementation. ACM Computing Surveys (Csur), 51(4), pp.1-35.
[17] Abadi, M., Chu, A., Goodfellow, I., McMahan, H.B., Mironov, I., Talwar, K. and Zhang, L., 2016, October. Deep learning with differential privacy. In Proceedings of the 2016 ACM SIGSAC conference on computer and communications security (pp. 308-318).
[18] Xiang, X., Wang, S., Huang, H., Qian, Y. and Yu, K., 2019, November. Margin matters: Towards more discriminative deep neural network embeddings for speaker recognition. In 2019 Asia-Pacific Signal and Information Processing Association Annual Summit and Conference (APSIPA ASC) (pp. 1652-1656). IEEE.
[19] Shokri, R. and Shmatikov, V., 2015, October. Privacy-preserving deep learning. In Proceedings of the 22nd ACM SIGSAC conference on computer and communications security (pp. 1310-1321).
[20] Bonawitz, K., Eichner, H., Grieskamp, W., Huba, D., Ingerman, A., Ivanov, V., Kiddon, C., Konečný, J., Mazzocchi, S., McMahan, B. and Van Overveldt, T., 2019. Towards federated learning at scale: System design. Proceedings of machine learning and systems, 1, pp.374-388.
[21] Smith, V., Chiang, C.K., Sanjabi, M. and Talwalkar, A.S., 2017. Federated multi-task learning. Advances in neural information processing systems, 30.
[22] Ogburn, Monique, Claude Turner, and Pushkar Dahal. "Homomorphic encryption." Procedia Computer Science 20 (2013): 502-509.
[23] Hegiste, Vinit & Legler, Tatjana & Ruskowski, Martin. (2022). Application of federated learning in manufacturing. 10.48550/arXiv.2208.04664.