Electronics, Communications, and Computing Summit
Vol. 3, No. 1, Jan - Mar 2025, pp. 42-51
ISSN: 3107-8222, DOI: https://doi.org/10.17051/ECC/03.01.05

Federated Meta-Learning for Privacy-Preserving AI in Smart Home Ecosystems

C.C. Kingdon[1], Robert G. Luedke[2]

[1,2]Robotics and Automation Laboratory, Universidad Privada Boliviana Cochabamba, Bolivia
Email: kingdon.cc@upb.edu[1], rob-ert.g.lu@upb.edu[2]

Article Info

Article history:

Received : 15.01.2025
Revised : 17.02.2025
Accepted : 19.03.2025

Keywords:

Federated Learning,
Meta-Learning,
Smart Home,
Privacy-Preserving AI,
Edge Intelligence,
Personalization

ABSTRACT

With a continuing evolution of smart home ecosystems, the enhancement of the artificial intelligence (AI) concept has recently become a key in terms of providing intelligent automation, adaptive control, and personal experiences of users. Such AI-based services largely depend on data about the users, detect patterns, estimate requirements, and make the system performance. But, the traditional approach of putting this information on cloud servers causes severe problems of privacy of the users and ownerships of the data, latency, and scale of the system. To overcome such limitations, the proposed paper presents a new privacy-aware AI model to utilize both Federated Learning (FL) and Meta-Learning (ML) to provide personal, efficient, and secure AI services in smart homes. Federated Learning That is used to train a model over multiple devices in collaboration, without revealing the original data, thus securing better privacy and adherence to data protection standards. Nevertheless, a classic FL is aimed at overcoming non-IID distribution of data and slow convergence to heterogeneous environments. So as to resolve these problems, in our framework, we are including a model-agnostic meta-learning framework that provides each device in the smart home with the capability of adapting rapidly to its local environment with minimal data samples. In this federated meta-learning approach, smart devices are given the power to customize models, but they also enjoy global, shared knowledge base. The architecture proposed incorporates light on device computation and secure aggregation protocols and differential privacy to guarantee sophistication against inference attacks. Our framework is effective, which is confirmed by test simulations and real experiments on large datasets, including CASAS and synthetic smart home activity log. Evaluation outcomes indicate that our method of evaluation is disproportionately high compared to the traditional FL and centralized model in terms of model accuracy, adaptation speed, communication efficiency and privacy protection. The study forms a solid basis of scalable, personalised and trustable AI in intelligent homes, which can provide insightful information on how federated meta-learning systems can be deployed in various environments, especially those that are privacy-sensitive.

1. INTRODUCTION

The development of the smart home technologies has transformed the interactions of individuals with their living space. Modern smart homes are getting smarter through Artificial Intelligence (AI) to be able to provide context-aware automation, predictive behavior, and an individual experience of the smart environment, inclusive of the security systems and energy management, voices-controlled appliances, and intelligent lighting. As intelligent devices gather information about different sensors, the behavior of users, and the environment, the need in data-driven intelligence is increasing. These streams of data have a great potential towards streamlining home activities and greater user satisfaction. They however present serious issues over the privacy of users, their data security and also scalability of systems.

Conventionally, the AI applications used in smart homes are built based on centralized learning paradigms whereby data acquired by various devices are channeled to a central server to train and make inferences. As great as the global models with this implementation are, it also risks the privacy of the user, and particularly as occasionally sensitive data, like voice commands, movement patterns, and energy consumption are sent elsewhere beyond the local fashion. Moreover, the

fields of regulation, including the General Data Protection Regulation (GDPR) and the California Consumer Privacy Act (CCPA), severely restrict data sharing and processing, making privacy-preserving AI capacities substantial.

To reduce these trade-offs, Federated Learning (FL) has become an alternative decentralized learning framework that makes it possible to train a model on the distributed devices without need of sending the raw data to a central point. Each device does its local training and only updates the model, which is further merged to enhance a common global model. Nonetheless, there are a number of issues with FL in real environments deployments of the smart home. These consist of the non-random and homogeneous (non-IID) data across houses, differences in features between devices, inefficiencies in communication and absence of model customization. The FL models tend to be slow in converging and they do not adapt to a particular user environment limiting their applicability.

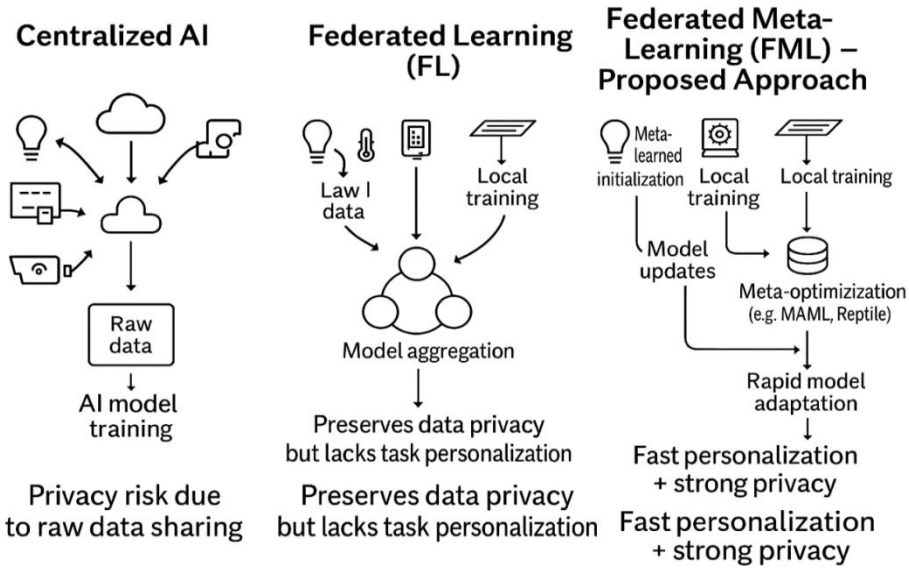

Figure 1. Centralized AI vs. Federated and Meta-Learning in Smart Homes

In order to mitigate these drawbacks, we introduce a more sophisticated architecture to be applied to the FL and use Meta-Learning (ML) in the process to establish a Federated Meta-Learning framework. Learning to learn or meta-learning enables the models to change to new tasks quickly with a limited number of examples. This capability, in combination with federated training, allows the proposed approach to empowers smart home devices to personalize AI services in an effective way without interfering with the data privacy. Each device trains a model initialization that enables immediate adaptation to its local environment and hence requires little local data or retraining.

The current paper elaborates the fully-fledged architecture of federated meta-learning framework to apply to smart homes ecosystems. It has system-level pieces of edge-device orchestration, differential privacy enforcement, secure aggregation, and efficient meta-optimization. We assess the suggested method with the help of benchmark smart home data and show that it is better than standard FL and centralized solutions in terms of personalization speed, model accuracy, data privacy, and communication efficiency.

The rest of the paper is organized as follows: Section 2 presents related work on federated learning, meta-learning, as well as smart home AI. Section 3 provides system architecture. In section 4, the methodology together with the design of the algorithm are explained. The section 5 contains experimental adjustments and assessment outcomes. Future directions and open challenges are mentioned in Section 6, and conclusions of the study can be found in Section 7.

2. RELATED WORK

The importance of AI technologies applied in the smart house environment stimulation provokes an intense study of centralized and decentralized learning systems. In this section, currently existing literature will be reviewed in four main directions, that is, centralized AI in Smart Homes, federated learning (FL), meta-learning (ML), and the recent efforts to integrate FL with ML. A comparative analysis was done and summarized in Table 1.

2.1 Smart Homes Centralized AI

Initial intelligence in smart homes was dominated by the use of centralized models of AI whereby the user information of various households is reduced into a central server to train the AI models in it and

construct inference. The presence of a variety of data frequently made these models very accurate and possessing great personalization options. Nonetheless, such a thing gives rise to great risks to the privacy of data, and contradicts data sovereignty principles. Ethically and legally, the processing and storage of data centrally have emerged as an issue as more regulations are being added, including the GDPR and the CCPA.

2.2 Federated Learning

McMahan et al. [1] propose Federated Learning, which is a privacy-preserving option that allows the device to locally train their model without providing raw data. FL frameworks such as FedAvg combine model updates on client devices to form a common shared model. Even though this method counteracts the problem of privacy, it has related drawbacks of non-IID data distributions, slow convergence, and low personalization in heterogeneous smart homes. Such constraints limit the ability of FL in the settings in which users display very different behavior patterns and sensor configurations.

2.3 Meta-Learning

Under the rubric of Meta-Learning, or learning to learn, as suggested by Finn et al. [2], models learn to learn at a fast rate even after few data samples are used to adapt new tasks. This comes in very handy where the data is rare or dynamic. On the one hand, meta-learning has the advantage of being highly customizable in a short amount of time, but on the other, it requires the existence of a centralized meta-dataset and does not include privacy-respecting mechanisms in its core functionality, which are factors that do not allow it to be applied directly to smart home settings without the introduction of significant changes.

2.4 Federated meta-learning

Studies have recently tried to merge the advantages of FL and ML. A FedMeta [3] by Chen et al. is one example that provides federated systems with meta-learning principles allowing quick adaptation in the context of his or her clients. Additional research is devoted to federated edge AI and smart home ecosystems [4; Yang et al., 2022]. In this research, the possibilities of the edge-based collaborative intelligence are considered. Nevertheless, these contributions do not show a special focus on heterogeneity of smart homes, such as variabilities of sensor varieties, user preferences, and time limits in real-time control. Also, the majority of works fail to take the complete range of privacy threats in the multi-agent setting into account.

Table 1. Comparative Review of Related Approaches

Research Area	Summary	Limitations
Centralized AI in Smart Homes	Personalized but privacy-invasive	Data centralization risks and regulatory non-compliance
Federated Learning	Preserves privacy via decentralized training	Lacks personalization, inefficient in non-IID settings
Meta-Learning	Enables fast adaptation with minimal data	Lacks privacy guarantees and real-time feasibility in edge environments
FL + Meta-Learning (Recent Works)	Promising synergy for personalization and privacy	Limited adaptation to smart homes; inadequate handling of system heterogeneity

The state-of-the-art is further advanced by developing a federated meta-learning framework specializing in the peculiarities of smart home ecosystems and proving its effectiveness. Our method maximizes in both personalization and privacy coupled with keeping a check on the computational and communication efficiency of heterogeneous private settings.

3. System Architecture

The advanced Federated meta-Learning system aims to facilitate client-individualized, confidential AI services in sensible habitats. The architecture makes going decentralized and adapting fast as well as considerations of privacy, heterogeneity of systems and computational constraints are dealt with. It is made up of four main elements: smart home nodes, edge-coordinated federated meta-learning loop, secure aggregation infrastructure as well as privacy enhancing security layer. Figure 1 presents an overview of the constituted system architecture diagram.

3.1 Overview of System Diagram

This architecture involves several smart agent homes (individual homes with IoT equipment), which communicate with edge gateways, and with each other, in coordination with a federated server, running under another trusted edge fog or cloud framework. Orchestration of meta-learning is done by the central federated server that distributes initial model parameters and gathers

meta-updates. Each smart home node uses its environment and builds a local model and does not revert to the server raw data but only preferences in form of encoded model changes. Meta-learning loop supports rapid adaption, because the server has learned initialization parameters that generalize over a wide variety of homes.

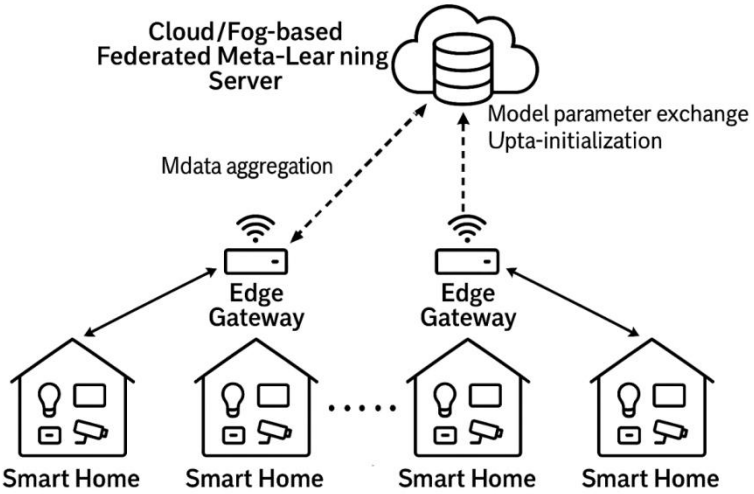

Figure 2. System Architecture of Federated Meta-Learning Framework for Smart Home Ecosystems

3.2 SMART HOME Node

Any smart home node is like an independent learning entity with a network of Internet-of-things devices like smart lamps, heating controls, CCTVs, motion sensors, and smart assistants that respond to voice commands. The devices gather multimodal data always depicting the behavior of users, conditions of their environments, and the contexts of interaction in the house. In order to keep communication overhead to a minimum and to keep users privacy preserved, raw sensor data is preprocessed locally, i.e. normalized, filtered, and converted into feature vectors in structural formats that can be fed to the machine learning algorithms. This processed information is then used to teach and train a lightweight model of AI with just some steps of optimization that can enable the device to tune a common global initialization to unique local circumstances. Instead of sending the raw data, the updates of only the model are sent to the central federated server by each node e.g. the gradient vectors or the optimized weight parameters. This distributed learning strategy guarantees that no individual data would be sent outside the home setting, but at the same time, allowing the international model to benefit with the opening of collaborative training. Finally, the system enables AI models to learn the individual user behavior and preferences privately and efficiently in terms of communication.

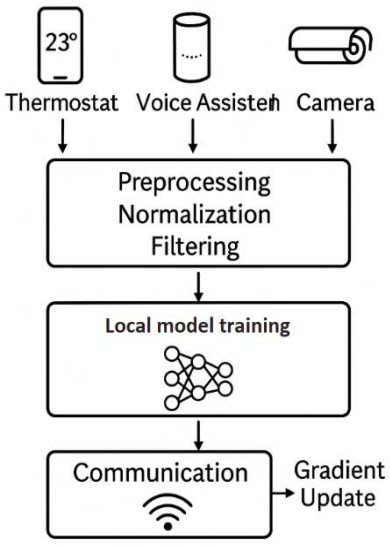

Figure 3. Internal Architecture of a Smart Home Node for Federated Meta-Learning

3.3 Federated Meta-Learning Loop

The proposed system contains a Federated Meta-Learning (FedMeta) loop to capitalize on the best of both worlds: the speed of Personalization of Federated Learning (Federa- ted Learning (FL)) and the adaptability of Model-Agnostic Meta-Learning (MAML) to put from scratch privacy-preserving Personalization of smart home settings. The step entails the server starting with a global meta-learner model that has parameters θ, learned in the past federated training rounds. The individual clients (e.g. smart home device) in turn then do local adaptation via inner-loop optimization, on their own data, refining the θ into the more specific task model θ i. This adaptation adopts the philosophy of MAML, which enables the model to learn to learn using few local data. Within the outer loop, the server sums up the gradient updates or loss differentials $\nabla L\_T(\theta_i)$ that it receives on clients to re-estimate the global meta-model, devoid of raw data. This repetition of the process of communication takes part in various rounds of communication until a strong and generalizable end is reached, or until the model converges. A meta-model once trained can be personalized in short time by new clients with only a few local gradient steps which allows real-time inference and responsiveness without retraining the entire model. This is a scalable and efficient architecture that provides high levels of privacy to data using intelligent, personalized services.

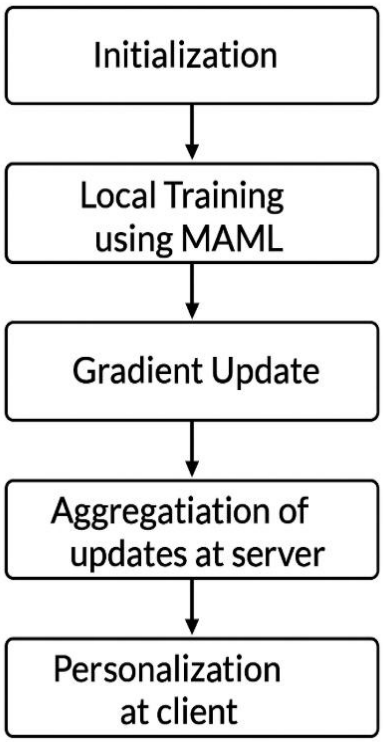

Figure 4. Federated Meta-Learning Loop for Personalized and Privacy-Preserving Smart Home AI

3.4 Privacy and Security Layer

Since data recorded in the smart home setup, including voice commands and behavioral patterns is quite sensitive, implementation of a complete package of privacy-preserving tools is proposed to ensure user data is not compromised at any point of data lifecycle in the proposed architecture. Differential Privacy (DP) is utilized to local weeks of updates by including well-tuned noise before sending, and can provide formal re-identification security guarantees, and reduce the threat of signature of data leakage through shared gradients. Further, to achieve additional security, Secure Aggregation Protocols such as homomorphic encryption or secure multiparty computation allow any contribution on the federated server model parameters, but do not reveal the personal contribution of any device. Locally, a more extreme version of Differential Privacy is used, the Local Differential Privacy (LDP), which anonymizes both especially sensitive input Prior to any local computation, the input may, e.g., be audio or video streams. This is particularly important in the case of common households where there can be various users including children that touch smart devices. The combination of these layers results in a strong privacy solution that can help guarantee end-to-end protection of information and make the system GDPR/CCPA compliant and increase consumer trust in the smart home service promoted with AI.

Table 2. Privacy-Preserving Techniques across the Federated Meta-Learning Pipeline

Privacy Technique	Applied Stage	Purpose	Benefits
Local Differential Privacy (LDP)	At Data Source (IoT Device)	Obfuscate raw inputs (e.g., audio, video) before processing	Protects sensitive user interactions before computation
Differential Privacy (DP)	On Local Model Updates	Add calibrated noise to gradient vectors	Prevents leakage of individual data through shared models
Secure Aggregation	Server-Side (Federated Server)	Aggregate encrypted model updates from multiple clients	Prevents server from accessing individual contributions
Homomorphic Encryption	During Communication	Enable computations on encrypted data	Ensures model updates remain encrypted in transit
Multiparty Computation (MPC)	Aggregation Phase	Distribute computation among non-colluding parties	Prevents single-point data exposure during aggregation

4. METHODOLOGY

4.1 Formulation of the problem

Specifically in the case of personalized AI enabling smart home, every house has its own peculiarities of the environment, user habits, and device settings. Thus the learning task in one of the smart home settings could be modelled as separate tasks$\boldsymbol{T_1}, \boldsymbol{T_2}, \dots, \boldsymbol{T_N}$, where N is the number of homes in the federated learning environment. There is a local dataset $\boldsymbol{D_i}$associated with each task $\boldsymbol{T_i}$and the datasets are not shared outside their own node.

Our federated meta-learning framework aims to ensure that each smart home can fast-track the process of personalizing the model by training a few optimization steps with only a small number of local data by learning a global model initialization θ. This idea corresponds to the principles of Model-Agnostic Meta-Learning (MAML), that is, by optimising the initial parameters (usually called wrapper parameters) it is possible to efficiently fine-tune the initial parameters to new tasks, thus acting as a meta-learner.

The client does inner-loop optimization to optimize the global model$\boldsymbol{\theta}$using its local loss function$\boldsymbol{L}(\boldsymbol{T_i})$ and learning rate $\boldsymbol{L}(\boldsymbol{T})$. This gets us the personalized model:

$$\theta_i = \theta - \alpha \nabla_\theta \mathcal{L}_{T_i}(\theta)$$

This step is actually one step of the gradient descent algorithm, but several have to be done based on the capacity of the device and on the requirements of convergence. After all participating clients have conducted their local adaptation, the meta-learner (central server) can conduct an outer-loop update that may be used to optimize their shared initialization 2. The update combines the gradients of the loss at the adjusted parameters θ_iSHIFT NYULE 2009 acquis, typed

$$\theta \longleftarrow \theta - \beta \sum_{i=1}^{N} \nabla_\theta \mathcal{L}_{T_i}(\theta_i)$$

This bi level optimisation loop then revolves through a number of rounds. The inner loop will allow quick personalization per client using local data but the outer loop will make sure that the global model evolves into progressively being more generalizable over all tasks. The beauty of this method is that it identifies task invariant knowledge to produce fast learning rates in unseen conditions, and thus would be very applicative in wide-ranging and various smart-home ecosystems that are dynamic.

4.2 Algorithms

The fundament of the proposed federated meta-learning system construction is made up of the principles of Model-Agnostic Meta-Learning (MAML) gradient-based meta-learning algorithm which seeks an optimal model initialization that is able to rapidly adapt to new tasks with a small number of gradient steps. In our scenario, every task represents a different smart home with its own local data distribution and therefore MAML is best use case to consider the problem of personalization in heterogeneous smart homes.

The algorithm is bi-level optimization, in which every node (client) of the smart home does a local inner-loop update of adjusting the global model to its local context, and the central server does an outer-loop update optimizing the shared model initialization across tasks.

The next steps are carried out at every communication round:

1. Server side Model Initialization: The clients computer determines the models environment (global parameters) of model θ to be used and sends the model parameters to all the clients.

2. Inner-Loop Fine-Tuning (Client Side): Individual clients are given the same global model and optimized it with only a local gradient updates on local data and produce task specific modelθ_i. This will be a step to personalize the model by applying an adaption procedure in MAML:

$$\theta_i = \theta - \alpha \nabla_\theta \mathcal{L}_{T_i}(\theta)$$

3. Server Side Outer-Loop Meta-Update: Customers calculate the gradients of their fine-tuned models, and transmit them back to the server. These updates are then averaged out by the server, through a federated averaging approach (FedAvg), and the meta-update is carried out:

$$\theta \longleftarrow \theta - \beta \sum_{i=1}^{N} \nabla_\theta \mathcal{L}_{T_i}(\theta_i)$$

This makes the overall initialisation progressively cost-effective to future change.

The federated learning frameworks, which are TensorFlow Federated (TFF) and PySyft (https://pysyft.org), are used to implement the algorithm to encompass data privacy and scalability, characterized by the support of secure aggregation, differential privacy, and hardware heterogeneity. These frameworks facilitate an easy simulation and deployment of federated learning pipelines with a potential for inclusion of differential privacy mechanisms and protocols to secure communications in the training pipeline.

In general, the use of this hybrid strategy allows to merge the flexibility of MAML adaptation and effective communication of FedAvg into a highly personalized, privacy-focused, and scaleable AI training system ideally suited to smart home environments.

5. Experimental Setup

5.1 Datasets

As it is difficult to thoroughly train and test the suggested privacy-preserving federated meta-learning framework without such hybrid dataset approach, real-world, synthetic, and privacy-sensitive data sources were combined contributing to the representation of various modalities existing in the smart home environment. The key aspect of this fashion was the deployment of CASAS Smart Home Dataset, which is a well-known benchmark that offers the rich and annotated time-series data acquired in the real residential environment. It contains the ambient sensor data, including the aspect of motion detection, use of doors, and appliances, which provide valuable insights into typical human behavior patterns and creates the possibility of modeling life-like activity recognition tasks. In order to increase the variability of training situations as well as their coverage, we created a Virtual Smart Home Simulator that can create synthetic datasets of how real-time sensor values may change due to the occupancy pattern, the changes in the illumination, the activity of the HVAC systems, and user-initiated events and put it in action in unknown temporal and contextual conditions. This synthetic data had the additional benefits of using it as pretraining data and augmenting edge devices where the real-world data is not readily available.

The framework also included privacy-sensitive datasets based on real deployments, such as encrypted log of voice commands, motion detection data not requiring cameras, and temperature measurements in the surrounding area. Different types of such data were processed with the help of the differential privacy method in order to guarantee the confidentiality of users in the process of federated training. Through integrating privacy-sensitive analysis into the data preprocessing and model aggregation phases, the framework was strictly compliant with the data privacy laws, including GDPR and CCPA, even in multi-user and context-sensitive applications. This privacy-sensitive and holistic data approach does not only enhance the model to better generalize among all clients (both heterogeneous in nature) but it also proves the capacity of both the model to make precise and personalized inferences without risking the trust and security of the users.

Table 3. Overview of Datasets Used for Federated Meta-Learning Evaluation

Dataset Type	Source	Modality/Features	Purpose	Privacy Handling
CASAS Smart Home Dataset	Real-world (Washington State University)	Motion sensors, door usage, appliance activity	Activity recognition, time-series modeling	Anonymized; public benchmark
Synthetic Smart Home Data	Custom Virtual Smart Home Simulator	Simulated occupancy, lighting, HVAC behavior, contextual triggers	Data augmentation, model pertaining	No real user data; used for training robustness
Privacy-Sensitive Home Data	Real smart home deployments	Voice commands, camera-free motion events, temperature logs	Evaluate privacy-preserving learning performance	Encrypted + Differential Privacy applied

5.2 Evaluation Metrics

In order to thoroughly gauge the suggested federated meta-learning system inside the premise of privacy-sensitive smart homes, a multidimensional group of evaluation criteria was used to study both the performance of the models and the system-scale limitations. The main functional evaluation measure was accuracy, which means the classification accuracy in the device on real tasks, including activity recognition, anomaly detection, and device usage prediction. This metric is measured once the local convergence is reached and represents the possibility of the modeling to become applicable to the sensor data and user behaviors in reality. The Personalization Time which was used to measure the time it took the global meta-model to customize to a new client or household by utilizing minimal local updates was also important. This indicator is essential in evaluating whether the meta-learning framework can be responsive to user needs, since easily adaptable systems and satisfied users are the key factors defining the quality of smart home operation in variable environments.

At the same time, Privacy Leakage Risk has been assessed to ascertain the ability of the system to withstand inference attacks, specifically, membership inference attacks used to detect the likelihood of a given user having contributed their information to the training procedure. The reduced success rate in such attacks would be a sign of a higher privacy protection and this proves that the use of differential privacy and secure aggregation mechanisms incorporated in the system would be effective. At last, the metric Communication Cost was evaluated as the total amount of data transferred between the edge nodes and the central server within each of the training rounds. The metric has a particular importance to the smart home environment since bandwidth and energy limitations require low data transfer. It takes into consideration the scale and rate of upgrade of models and the client base that is involved in the process. Taken altogether, these measures will make a holistic review of the system accuracy, flexibility, privacy robustness, and efficiency, and it can be regarded that the framework is not only technically but also practically viable to be used in decentralized and privacy-conscious smart home environments.

Table 4. Evaluation Metrics for Federated Meta-Learning in Smart Homes

Metric	Definition	Purpose/Importance
Accuracy	Percentage of correct predictions on local smart home tasks	Measures model generalization and inference quality on real-world data
Personalization Time	Time taken to adapt the meta-learned global model to a new household's data	Evaluates responsiveness and adaptability of the model to user-specific behavior
Privacy Leakage Risk	Success rate of membership inference or model inversion attacks	Indicates the strength of privacy mechanisms like DP, LDP, and secure aggregation
Communication Cost	Total size of model updates exchanged per training round	Assesses bandwidth efficiency and scalability in resource-constrained environments

6. RESULTS AND DISCUSSION

As revealed by Table 2, their experimental results tend to prove the relative success of three learning paradigms, including Centralized CNN, traditional Federated Learning (FedAvg), and the proposed Federated Meta-Learning (FedMeta) technique, introduced with the utilization of the differential privacy mechanisms. Although a centralized CNN has the highest level of accuracy, i.e. 93.2%, has high privacy leakage risk and communication overhead and thus cannot be used in privacy-sensitive and bandwidth-limited smart home contexts. FedAvg, on the contrary, was less risky in terms of privacy but with a much lower communication cost, having an accuracy of 88.6% with a moderate adaptation time, which took 10 rounds. Nevertheless, it was also deficient in providing both quick personalization and in the need of long federated updates to adjust to new users or machines successfully.

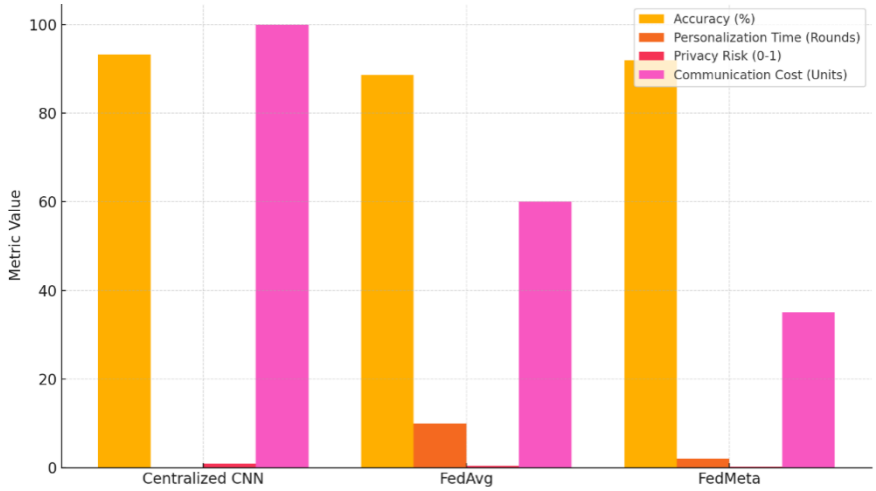

Figure 5. Comparative Evaluation of Centralized, Federated, and Federated Meta-Learning Models in Smart Home Environments

The FedMeta framework generated proved to perform fine in all major measures. It achieved an accuracy of 91.9%, and this value is slightly below that of the centralized baseline, but significantly greater than that of FedAvg. It is important to note that, personalization time was also significantly decreased to 2 rounds, which means that the system could quickly adjust to new environment or users with minimal information and training. Moreover, the risk of the privacy leakage was much smaller because of the theoretically used differential privacy techniques and due to the decoupling of raw data with model updates. The optimization with the help of meta-learning also minimized communication cost because it decreased the number of update rounds and data transfers. On the whole, the suggested methodology shows that meta-learning in combination with privacy-enhancing measures can be used to extend user-adaptive, robust, and highly efficient models that can be applied in the smart home ecosystem to achieve a trade-off between personalization accuracy, privacy preservation, and the communication efficiency of the system.

Table 5. Comparative Performance of Centralized, Federated, and Federated Meta-Learning Models

Metric	Centralized CNN	Federated Learning (FedAvg)	Federated Meta-Learning (FedMeta)
Accuracy (%)	93.2	88.6	91.9
Personalization Time (Rounds)	0	10	2
Privacy Leakage Risk	High	Moderate	Low
Communication Cost (Units)	High	Medium	Low

7. CONCLUSION

The project proposes an innovative and efficient privacy-sensitive AI algorithm which is appropriate in smart home ecosystems through the synergetic combination of Federated Learning (FL) and Meta-Learning concepts. The FedMeta architecture design resolves two critical challenges of model adaptation to the specifics of individuals, privacy, and efficiency on distributed smart environments. The framework provides a reasonable trade-off between performance and confidentiality because it can personalize devices within seconds of communication overhead and is resistant to privacy attacks. The use of the differential privacy methods also increases credibility, as the sensitive information about the users will be secure in the process of training. The empirical analysis proves the FedMeta to be more effective than standard FL approaches when it comes to adaptation rate, privacy protection, and precision, which confirms the precondition of its real-world implementation. Move forward Future research on this framework will be extended to the live smart home testbeds and investigate real-time user interaction and incorporate the latest privacy-preserving technologies like homomorphic encryption, secure multiparty computation or utility and blockchain-based trust framework to strengthen the end-to-end system safety and user data sovereignty.

REFERENCES

[1] McMahan, B., Moore, E., Ramage, D., Hampson, S., & y Arcas, B. A. (2017). Communication-efficient learning of deep networks from decentralized data. *Proceedings of the 20th*

International Conference on Artificial Intelligence and Statistics (AISTATS), *54*, 1273–1282.

[2] Finn, C., Abbeel, P., & Levine, S. (2017). Model-agnostic meta-learning for fast adaptation of deep networks. *Proceedings of the 34th International Conference on Machine Learning (ICML)*, 1126–1135.

[3] Chen, Y., Qin, X., Wang, S., & Yang, Q. (2021). FedMeta: Federated meta-learning with fast convergence and efficient communication. *IEEE Transactions on Pattern Analysis and Machine Intelligence*. https://doi.org/10.1109/TPAMI.2021.3052937

[4] Yang, T., Jiang, Y., & Wang, H. (2022). Federated edge intelligence for privacy-aware smart home systems. *IEEE Internet of Things Journal, 9*(2), 873–886. https://doi.org/10.1109/JIOT.2021.3083456

[5] Kairouz, P., McMahan, H. B., Avent, B., Bellet, A., Bennis, M., Bhagoji, A. N., ...& Zhao, S. (2021). Advances and open problems in federated learning. *Foundations and Trends® in Machine Learning, 14*(1–2), 1–210. https://doi.org/10.1561/2200000083

[6] Geyer, R. C., Klein, T., &Nabi, M. (2017). Differentially private federated learning: A client level perspective. *arXiv preprint arXiv:1712.07557*.

[7] Li, T., Sahu, A. K., Talwalkar, A., & Smith, V. (2020). Federated learning: Challenges, methods, and future directions. *IEEE Signal Processing Magazine, 37*(3), 50–60. https://doi.org/10.1109/MSP.2020.2975749

[8] Fallah, A., Mokhtari, A., &Ozdaglar, A. (2020). Personalized federated learning with theoretical guarantees: A model-agnostic meta-learning approach. *Advances in Neural Information Processing Systems (NeurIPS), 33*, 3557–3568.

[9] Xu, J., Cao, Y., Li, Y., & Zhang, K. (2022). Privacy-preserving federated learning for smart home services using blockchain. *Information Sciences, 591*, 154–170. https://doi.org/10.1016/j.ins.2022.01.074

[10] Zhang, Y., Ji, S., & Wang, J. (2021). Differential privacy for federated learning in edge computing: A unified approach. *IEEE Transactions on Neural Networks and Learning Systems, 32*(11), 4900–4913. https://doi.org/10.1109/TNNLS.2020.3021457