

FORM 2
THE PATENT ACT, 1970
(39 OF 1970)
&
THE PATENT RULES, 2003
COMPLETE SPECIFICATION
[SEE SECTION 10 AND RULE 13]

TITLE: Advanced Federated Learning Architecture for Enhanced Privacy and Collaborative Data Training Across Heterogeneous Devices	
APPLICANTS:	
Dr. A. Raji Reddy	Professor, Mechanical Engineering, CMR Technical Campus KANDLAKOYA, MEDCHAL ROAD, HYDERABAD, TELANGANA, INDIA, 501401
Dr Vinoda Reddy	Assoc. Prof., Computer Science and Engineering (AI & ML), CMR Technical Campus KANDLAKOYA VILLAGE, MEDCHAL MANDAL, R. R DISTRICT, HYDERABAD 501401 TELANGANA, INDIA

**The following specification particularly describes the invention and the manner
in which it is to be performed**

**ADVANCED FEDERATED LEARNING ARCHITECTURE FOR
ENHANCED PRIVACY AND COLLABORATIVE DATA TRAINING
ACROSS HETEROGENEOUS DEVICES**

FIELD OF THE INVENTION

[001] Various embodiments of the present invention generally relate to federated learning architecture. More particularly, the invention relates to an advanced federated learning architecture for enhanced privacy and collaborative data training across heterogeneous devices.

BACKGROUND OF THE INVENTION

[002] The rapid advancement of machine learning technologies has significantly transformed various industries, driving innovations in fields such as healthcare, finance, and smart devices. Federated learning has emerged as a pivotal approach to addressing some of the critical challenges associated with centralized machine learning, particularly concerning data privacy and the management of distributed data sources.

[003] In traditional machine learning paradigms, data is collected from multiple sources and centralized in a single location for model training. While this approach allows for the aggregation of extensive data sets, it raises substantial concerns regarding data privacy and security. Centralized data storage poses risks of unauthorized access, data breaches, and potential misuse of sensitive information. These concerns are particularly acute when dealing with personal, financial, or medical data, where stringent privacy regulations are in place.

[004] Federated learning offers a promising alternative by enabling collaborative training of machine learning models across distributed devices while keeping data localized. Instead of transferring raw data to a central server, federated learning processes data on individual devices and only shares model updates, which helps mitigate privacy risks. Despite its advantages, federated learning introduces its own set of challenges, including:

Privacy Preservation: Ensuring that model updates do not inadvertently expose sensitive information is crucial. Techniques such as encryption and differential privacy must be integrated to protect data during training and communication.

Heterogeneity of Devices: Federated learning involves a variety of devices with different hardware capabilities, software environments, and data types. This heterogeneity can complicate the coordination and integration of model updates from diverse sources.

Secure Aggregation: Combining model updates from multiple devices without revealing individual contributions requires sophisticated aggregation techniques. Ensuring that the aggregation process does not compromise privacy is a key challenge.

Efficient Communication: Managing the communication between numerous devices and a central server efficiently and securely is essential for maintaining system performance and data integrity.

[005] The invention addresses these challenges by providing an advanced federated learning architecture that enhances privacy protection, supports a diverse range of devices, and employs secure aggregation and communication protocols. This architecture allows for effective collaborative data training while preserving the

confidentiality of sensitive information and ensuring compliance with data protection regulations.

[006] By integrating privacy-preserving mechanisms, secure communication protocols, and scalable aggregation techniques, the invention represents a significant advancement in the field of federated learning. It aims to bridge the gap between effective machine learning and stringent privacy requirements, making it a valuable solution for a wide range of applications where data privacy and security are paramount.

[007] One or more advantages of the prior art are overcome, and additional advantages are provided through the invention. Additional features are realized through the technique of the invention. Other embodiments and aspects of the disclosure are described in detail herein and are considered a part of the invention.

SUMMARY OF THE INVENTION

[008] The invention pertains to an advanced federated learning architecture designed to enhance privacy and facilitate collaborative data training across a diverse set of devices. This architecture includes a central server and a plurality of heterogeneous devices, each equipped to perform local machine learning tasks using its own data.

[009] The method involves initializing the federated learning system, distributing a global machine learning model to all devices, and conducting local training while applying privacy-preserving techniques. These techniques include encryption of local data and model updates, differential privacy, and secure aggregation methods such as homomorphic encryption and secure multi-party computation.

[010] Once local model updates are encrypted, they are transmitted to the central server, where they are aggregated to create an updated global model. This updated model is then redistributed to the devices for further training. The iterative process continues until the global model achieves the desired performance level.

[011] Throughout the process, secure communication protocols ensure data integrity and privacy. The architecture's design allows for effective collaborative learning while preserving individual data confidentiality, accommodating diverse devices, and complying with data protection regulations. This approach enhances model performance by integrating varied data sources and maintains robust privacy and security standards.

BRIEF DESCRIPTION OF THE FIGURES

[012] The accompanying figures where like reference numerals refer to identical or functionally similar elements throughout the separate views and which together with the detailed description below are incorporated in and form part of the specification, serve to further illustrate various embodiments and to explain various principles and advantages all in accordance with the invention.

[013] FIG. 1 is a diagram that illustrates an advanced federated learning architecture, in accordance with an embodiment of the invention.

[014] FIG. 2 is a diagram that illustrates a flow diagram with a method for enhanced privacy-preserving collaborative data training across heterogeneous devices using an advanced federated learning architecture, in accordance with an embodiment of the invention.

[015] Skilled artisans will appreciate the elements in the figures are illustrated for simplicity and clarity and have not necessarily been drawn to scale. For example, the dimensions of some of the elements in the figures may be exaggerated relative to other elements to help to improve understanding of embodiments of the present invention.

DETAILED DESCRIPTION OF THE INVENTION

[016] While various embodiments of the invention have been shown and described herein, it will be obvious to those skilled in the art that such embodiments are provided by way of example only. Numerous variations, changes, and substitutions may occur to those skilled in the art without departing from the invention. It should be understood that various alternatives to the embodiments of the invention described herein may be employed. It shall be understood that different aspects of the invention can be appreciated individually, collectively, or in combination with each other.

[017] FIG. 1 is a diagram that illustrates an advanced federated learning architecture (100), in accordance with an embodiment of the invention.

[018] Referring to FIG. 1, the system 100 the comprises a plurality of heterogeneous devices (102), a central server (104), a privacy-preserving mechanism (106), a collaborative data training module (108), an aggregation engine (110), and a communication protocol (112)

[019] The advanced federated learning architecture (100) is designed to enhance privacy and facilitate collaborative data training across a diverse range of devices. This architecture comprises several key components:

[020] Plurality of Heterogeneous Devices (102): The architecture includes a variety of devices (102), each equipped to perform local machine learning tasks using their respective local data. These devices can differ in their hardware, software, and data types, reflecting the real-world diversity of computing environments.

[021] The plurality of heterogeneous devices (102) refers to a diverse set of computing devices that participate in the federated learning process. These devices are characterized by their varied hardware, software, and data capabilities, reflecting a broad spectrum of real-world environments. The heterogeneity of these devices is a key feature of the architecture, as it allows for the integration of a wide range of data sources and computational resources.

[022] Device Diversity: The devices (102) can include, but are not limited to, smartphones, tablets, personal computers, servers, IoT devices, and edge computing devices. Each device may have different processing power, memory capacity, and

storage capabilities. This diversity ensures that the federated learning system can leverage a wide array of data and computational resources.

[023] Local Machine Learning Tasks: Each device (102) is configured to perform local machine learning tasks using its own local data. This involves training a local model on data that resides on the device itself. The local training process is designed to be resource-efficient, taking into account the device's computational limitations and energy constraints.

[024] Local Data: The local data used for training on each device (102) can vary significantly. It may include user-generated content, sensor data, or application-specific data. The heterogeneity of data sources enhances the robustness and generalizability of the global model, as it captures a wide range of data distributions and patterns.

[025] Privacy Considerations: Given the diverse nature of the devices (102) and their data, privacy considerations are paramount. Each device handles sensitive and potentially personal information. The federated learning architecture ensures that data privacy is maintained by processing data locally and transmitting only model updates, rather than raw data, to the central server (104).

[026] Adaptability and Scalability: The architecture is designed to accommodate an evolving set of devices (102) that may join or leave the network over time. The system is scalable, meaning that it can handle an increasing number of devices without significant performance degradation. This adaptability is crucial for maintaining the effectiveness of the federated learning process in dynamic environments.

[027] Interoperability: The devices (102) are expected to operate under a common framework but may use different software platforms and communication protocols. The system includes mechanisms to ensure interoperability and effective communication between these diverse devices, facilitating seamless integration and collaboration.

[028] Central Server (104): The central server (104) plays a crucial role in coordinating the federated learning process. It is responsible for aggregating model updates received from the heterogeneous devices (102) and ensuring that privacy-

preserving techniques are implemented throughout the process. The server orchestrates the learning algorithm, manages communication, and ensures that the global model is updated accurately.

[029] Privacy-Preserving Mechanism (106): Integrated into the architecture is a privacy-preserving mechanism (106) designed to safeguard sensitive information. This mechanism encrypts local data and/or model updates before they are transmitted to the central server (104). By preventing unauthorized access, it ensures the privacy and security of the data throughout the learning process.

[030] The privacy-preserving mechanism (106) is a critical component of the advanced federated learning architecture (100), designed to ensure the confidentiality and security of sensitive information throughout the federated learning process. This mechanism addresses the challenge of preserving privacy while enabling collaborative model training across heterogeneous devices (102). The privacy-preserving mechanism (106) incorporates various techniques and technologies to protect data at multiple stages:

[031] Encryption of Local Data:

[032] Data Encryption: Before any local data is used for training or transmitted, it is encrypted using robust encryption algorithms. This ensures that even if data is intercepted or accessed by unauthorized parties, it remains unreadable and secure.

[033] Homomorphic Encryption: In some implementations, homomorphic encryption is employed to allow computations to be performed on encrypted data without decrypting it. This enables the central server (104) to aggregate model updates securely while maintaining data confidentiality.

[034] Encryption of Model Updates:

[035] Update Encryption: During the training process, each device (102) generates model updates based on its local data. These updates are encrypted before being sent to the central server (104). This encryption protects the model updates from being accessed or tampered with during transmission.

[036] Secure Aggregation: The central server (104) uses secure aggregation techniques to combine encrypted model updates from multiple devices (102). This approach ensures that individual updates are not disclosed, maintaining the privacy of the local data used to generate them.

[037] Differential Privacy:

[038] Noise Addition: To further protect individual data points, differential privacy techniques are applied. This involves adding carefully calibrated noise to the local model updates or data before transmission. The noise obscures the contribution of any single data point, ensuring that the overall privacy of the data is preserved.

[039] Privacy Budget Management: Differential privacy implementations manage a privacy budget to control the amount of noise added and ensure that privacy guarantees are maintained while still allowing useful learning.

[040] Anonymization Techniques:

[041] Data Anonymization: Local data is anonymized before being used for training to prevent identification of individuals or sensitive information. This can involve removing or obfuscating personal identifiers and sensitive attributes.

[042] Pseudonymization: In some cases, pseudonymization is used to replace personal identifiers with pseudonyms, ensuring that the data remains anonymous while still allowing meaningful analysis.

[043] Access Control and Authentication:

[044] Access Control: The privacy-preserving mechanism (106) includes access control mechanisms to ensure that only authorized entities can access or process data. This prevents unauthorized access and potential breaches of data privacy.

[045] Secure Authentication: Strong authentication protocols are employed to verify the identity of devices (102) and users participating in the federated learning process. This ensures that only legitimate devices can join the network and contribute to the learning process.

[046] Compliance with Privacy Regulations:

[047] Regulatory Compliance: The privacy-preserving mechanism (106) is designed to comply with relevant data protection regulations and standards, such as GDPR, HIPAA, or CCPA. This ensures that the federated learning process adheres to legal and ethical standards for data privacy.

[048] Collaborative Data Training Module (108): The collaborative data training module (108) enables the secure exchange of model updates between the heterogeneous devices (102) and the central server (104). It facilitates the collaborative aspect of federated learning while maintaining data confidentiality, ensuring that individual data points remain private.

[049] Aggregation Engine (110): Located on the central server (104), the aggregation engine (110) is responsible for combining model updates from different devices (102). This process enhances the performance of the global model by incorporating diverse data while preserving the privacy of individual contributions. The engine ensures that the aggregation process does not compromise data confidentiality.

[050] The aggregation engine (110) is a pivotal component of the advanced federated learning architecture (100), located on the central server (104). Its primary function is to combine model updates received from the plurality of heterogeneous devices (102) into a unified global model. This process is performed with a strong emphasis on maintaining the privacy and security of the individual contributions. Here's a comprehensive overview of its key features and operations:

[051] Collection and Integration of Model Updates:

[052] Update Collection: The aggregation engine (110) receives encrypted model updates from the heterogeneous devices (102). These updates represent the results of local machine learning tasks performed on each device's data.

[053] Integration Process: The engine integrates these model updates to form a global model that reflects the combined knowledge and learning derived from all participating devices. This integration is performed in a manner that preserves the privacy of the individual updates.

[054] Secure Aggregation Techniques:

[055] Homomorphic Encryption: To ensure that model updates are combined without exposing individual updates, the aggregation engine (110) may use homomorphic encryption techniques. These techniques allow computations to be performed on encrypted data, enabling secure aggregation without decrypting the updates.

[056] Secure Multi-Party Computation (SMPC): Another approach employed may be secure multi-party computation, which allows the aggregation engine (110) to compute the combined model while keeping individual updates confidential.

[057] Privacy Preservation:

[058] Anonymization of Updates: During the aggregation process, the engine ensures that the model updates are anonymized to prevent any potential exposure of sensitive information related to specific devices (102) or their local data.

[059] Noise Addition: Differential privacy mechanisms may be integrated into the aggregation process to add noise to the aggregated model, further protecting against the risk of revealing sensitive data.

[060] Model Update Optimization:

[061] Weighted Aggregation: The aggregation engine (110) may implement weighted aggregation methods to account for the varying sizes and quality of data across different devices (102). This ensures that the contributions from each device are appropriately represented in the global model.

[062] Optimization Algorithms: Advanced optimization algorithms may be employed to efficiently combine the model updates and improve the overall performance of the global model. This can include techniques such as averaging, gradient-based methods, or other aggregation strategies.

[063] Version Control and Tracking:

[064] Model Versioning: The aggregation engine (110) maintains version control of the global model, tracking changes and updates made during each aggregation cycle. This ensures consistency and allows for the monitoring of model evolution over time.

[065] Tracking Updates: It also keeps track of the model updates received from each device, ensuring that all contributions are accounted for and that the aggregation process is transparent and traceable.

[066] Performance and Scalability:

[067] Efficiency: The aggregation engine (110) is designed to handle a large volume of model updates efficiently. It optimizes resource usage and computational power to manage the aggregation process effectively, even as the number of participating devices (102) increases.

[068] Scalability: The engine is scalable to accommodate the growth in the number of devices and data volume. It is capable of integrating updates from a growing network of devices without significant performance degradation.

[069] Interoperability and Compatibility:

[070] Support for Multiple Models: The aggregation engine (110) can support various machine learning models and algorithms, allowing it to integrate updates from different types of models and ensure compatibility across diverse devices (102).

[071] Flexible Integration: It is designed to work seamlessly with the communication protocol (112) and other components of the federated learning architecture (100), ensuring smooth integration and operation within the overall system.

[072] Communication Protocol (112): The communication protocol (112) supports secure and efficient data transfer between the heterogeneous devices (102) and the central server (104). It employs end-to-end encryption and secure authentication methods to protect data integrity and privacy during transmission. This protocol ensures that data remains secure from potential threats and unauthorized access.

[073] FIG. 2 is a diagram that illustrates a flow diagram 200 with a method for enhanced privacy-preserving collaborative data training across heterogeneous devices using an advanced federated learning architecture, in accordance with an embodiment of the invention.

[074] Step 202: The method for enhanced privacy-preserving collaborative data training begins by initializing a federated learning system with a plurality of

heterogeneous devices and a central server. Each device in the system is equipped with its own local data and machine learning capabilities, while the central server coordinates the overall federated learning process. This initialization involves setting up the necessary infrastructure for communication, ensuring that each device and the server are properly configured to participate in the collaborative learning process.

[075] Step 204: The next step involves distributing a global machine learning model to the plurality of devices. The global model, initially created or pre-trained on a central repository, is sent to each device. This model serves as the starting point for local training, allowing each device to build upon the existing knowledge while adapting the model to its specific data.

[076] Step 206: Each device then performs local machine learning training using its own local data. During this training process, privacy-preserving techniques are applied to ensure data confidentiality. These techniques may include data anonymization, differential privacy, or other methods to protect sensitive information. The local training allows each device to update the global model based on its unique data while ensuring that individual data points remain private and secure.

[077] Step 208: After completing local training, each device encrypts the local model updates before transmitting them to the central server. This encryption step is crucial for maintaining data privacy and security during transmission. The encrypted updates prevent unauthorized access and ensure that the information remains confidential as it travels across potentially insecure networks.

[078] Step 210: At the central server, the encrypted model updates from each device are aggregated to generate an updated global model. The aggregation process is designed to combine the contributions from all devices while preserving the privacy of the individual data used to generate the updates. Techniques such as homomorphic encryption or secure multi-party computation may be employed to ensure that the aggregation process does not compromise data confidentiality.

[079] Step 212: The updated global model is then distributed back to the plurality of devices for further local training. This iterative process allows each device to refine its

local model based on the updated global model, continuously improving the overall performance of the global model through collaborative learning.

[080] Step 214: Steps 206 through 212 are repeated iteratively until the global model reaches a desired performance level. This iterative training process ensures that the model benefits from contributions made by all devices, progressively enhancing its accuracy and effectiveness while maintaining privacy throughout each cycle.

[081] Step 216: Throughout the entire federated learning process, secure communication is maintained to protect data integrity and privacy. This includes using encryption and secure authentication methods for data transmission between devices and the central server. Ensuring secure communication is vital for preventing data breaches and preserving the confidentiality of the information shared during the learning process.

[082] Enhanced Privacy Protection: By employing privacy-preserving techniques such as encryption of local data and model updates, as well as differential privacy, the invention ensures that sensitive information remains confidential. The use of these techniques prevents unauthorized access and protects individual data from being exposed during the collaborative learning process.

[083] Collaborative Learning Across Diverse Devices: The architecture supports collaboration among a variety of heterogeneous devices. This diversity enhances the global model by incorporating data and learning from multiple sources, leading to a more robust and generalized model that performs better in real-world scenarios.

[084] Efficient Data Utilization: Local machine learning training allows each device to leverage its own data without requiring raw data to be transferred to a central location. This approach maximizes the utility of data while minimizing the need for extensive data transfer, reducing network bandwidth usage and associated costs.

[085] Scalability: The system is designed to scale efficiently with the number of participating devices. As more devices join the network, the federated learning process can accommodate increased data and model updates without significant performance degradation, making it suitable for large-scale applications.

[086] Secure Aggregation: The aggregation engine employs secure methods such as homomorphic encryption and secure multi-party computation to combine model updates without exposing individual contributions. This ensures that the central server can generate an updated global model while preserving the privacy of the local data used to create the updates.

[087] Iterative Improvement: The iterative nature of the method, where local training and model updates are repeatedly performed, allows for continuous improvement of the global model. This iterative approach ensures that the model evolves and adapts over time, incorporating new data and insights from all participating devices.

[088] Regulatory Compliance: The architecture is designed to comply with various data protection regulations and standards, such as GDPR, HIPAA, or CCPA. This compliance ensures that the federated learning process adheres to legal and ethical requirements for data privacy and security.

[089] Reduced Data Exposure: By only transmitting encrypted model updates rather than raw data, the invention minimizes the risk of data breaches and unauthorized data exposure. This approach maintains privacy while allowing for effective model training and enhancement.

[090] Improved Model Performance: The aggregation of diverse model updates from multiple devices contributes to a more accurate and effective global model. The combined knowledge from various data sources helps in addressing a wider range of scenarios and improving the overall performance of the model.

[091] Efficient Communication Protocol: The use of a secure and efficient communication protocol ensures that data transfer between devices and the central server is protected and optimized. This minimizes potential security risks and ensures reliable and fast communication throughout the federated learning process.

[092] Those skilled in the art will realize that the above-recognized advantages and other advantages described herein are merely exemplary and are not meant to be a complete rendering of all of the advantages of the various embodiments of the present invention.

[093] In the foregoing complete specification, specific embodiments of the present invention have been described. However, one of ordinary skill in the art appreciates that various modifications and changes can be made without departing from the scope of the present invention. Accordingly, the specification and the figures are to be regarded in an illustrative rather than a restrictive sense. All such modifications are intended to be included with the scope of the present invention and its various embodiments.

Dated 27th August 2024

Ganapathi S Naidu

IN/PA – 2312

Digitally Signed

Authorized Patent Agent for the Applicant

I/WE CLAIM:

1. An advanced federated learning architecture (100) comprising:
 - a plurality of heterogeneous devices (102), each device being configured to perform local machine learning tasks on local data;
 - a central server (104) configured to coordinate the federated learning process, aggregate updates from the plurality of devices, and ensure privacy-preserving techniques are applied;
 - a privacy-preserving mechanism (106) integrated into the architecture, capable of encrypting local data and/or model updates to prevent unauthorized access and ensure data privacy;
 - a collaborative data training module (108) that facilitates the exchange of model updates between the heterogeneous devices and the central server while maintaining data confidentiality;
 - an aggregation engine (110) on the central server (104) that combines model updates from different devices to improve the overall model performance while preserving individual data privacy;
 - a communication protocol (112) that supports secure and efficient data transfer between the heterogeneous devices and the central server.
2. The advanced federated learning architecture (100) of claim 1, wherein the privacy-preserving mechanism (106) includes a differential privacy module that adds noise to local model updates to further obscure sensitive information before transmission.
3. The advanced federated learning architecture (100) of claim 1, wherein the collaborative data training module (108) includes a data anonymization

component that transforms local data into an anonymized format before training to enhance data privacy.

4. The advanced federated learning architecture (100) of claim 1, wherein the aggregation engine (110) incorporates a secure multi-party computation technique to ensure that the aggregation of model updates is performed without revealing individual updates to the central server (104).
5. The advanced federated learning architecture (100) of claim 1, wherein the communication protocol (112) utilizes end-to-end encryption and secure authentication methods to protect data during transmission between the heterogeneous devices (102) and the central server (104).
6. A method for enhanced privacy-preserving collaborative data training across heterogeneous devices using an advanced federated learning architecture, the method comprising:

initializing a federated learning system with a plurality of heterogeneous devices and a central server;

distributing a global machine learning model to the plurality of devices;

performing local machine learning training on each device using local data while applying privacy-preserving techniques to ensure data confidentiality;

encrypting local model updates and transmitting them from each device to the central server;

aggregating the encrypted model updates at the central server to generate an updated global model while preserving the privacy of the local data;

distributing the updated global model back to the plurality of devices for further local training;

repeating steps iteratively until the global model reaches a desired performance level; and

ensuring secure communication throughout the process to maintain data integrity and privacy.

7. The method of claim 6, wherein step c) further includes applying a differential privacy technique to the local model training process, adding noise to the local model updates to obscure sensitive information before transmission to the central server (104).
8. The method of claim 6, wherein step c) includes anonymizing the local data using a data anonymization technique prior to performing local machine learning training, enhancing data privacy before any processing or transmission.
9. The method of claim 6, wherein step e) incorporates a secure multi-party computation approach to aggregate the encrypted model updates at the central server (104) without revealing individual updates or local data.
10. The method of claim 6, wherein step h) includes utilizing end-to-end encryption and secure authentication protocols during the secure communication process to protect data integrity and privacy throughout transmission between the heterogeneous devices (102) and the central server (104).

Dated 27th August 2024

Ganapathi S Naidu

IN/PA – 2312

Digitally Signed

Authorized Patent Agent for the Applicant

**ADVANCED FEDERATED LEARNING ARCHITECTURE FOR
ENHANCED PRIVACY AND COLLABORATIVE DATA TRAINING
ACROSS HETEROGENEOUS DEVICES**

ABSTRACT

The advanced federated learning architecture (100) provides a robust framework for privacy-preserving collaborative data training across heterogeneous devices (102). This architecture features a central server (104) that coordinates the federated learning process and aggregates updates from the devices. A privacy-preserving mechanism (106) ensures the confidentiality of local data and model updates through encryption. The system includes a collaborative data training module (108) to facilitate the secure exchange of model updates while maintaining data confidentiality. An aggregation engine (110) on the central server (104) combines these updates to enhance overall model performance while preserving individual data privacy. The communication protocol (112) supports secure and efficient data transfer between devices (102) and the central server (104), utilizing end-to-end encryption and secure authentication to safeguard data integrity throughout the process. This architecture addresses the need for secure, privacy-preserving collaborative machine learning across diverse devices.

FIG.1