Empowering Digital Agriculture: A Privacy-Preserving Framework for Data Sharing and Collaborative Research

Osama Zafar
Department of Computer and Data Sciences
Case Western Reserve University
Cleveland, OH 44106, USA
oxz23@case.edu

Rosemarie Santa González
School of Computer Science
Georgia Institute of Technology
Atlanta, GA 30332, USA
rosemarie.santa@gatech.edu

Mina Namazi
Department of Computer and Data Sciences
Case Western Reserve University
Cleveland, OH 44106, USA
mxn559@case.edu

Alfonso Morales
Department of Planning and Landscape Architecture
University of Wisconsin–Madison
Madison, WI 53706, USA
morales1@wisc.edu

Erman Ayday
Department of Computer and Data Sciences
Case Western Reserve University
Cleveland, OH 44106, USA
exa208@case.edu

Abstract—Data-driven agriculture, which integrates technology and data into agricultural practices, has the potential to improve crop yield, disease resilience, and long-term soil health. However, privacy concerns, such as adverse pricing, discrimination, and resource manipulation, deter farmers from sharing data, as it can be used against them. To address this barrier, we propose a privacy-preserving framework that enables secure data sharing and collaboration for research and development while mitigating privacy risks. The framework combines dimensionality reduction techniques (like Principal Component Analysis (PCA)) and differential privacy by introducing Laplacian noise to protect sensitive information. The proposed framework allows researchers to identify potential collaborators for a target farmer and train personalized machine learning models either on the data of identified collaborators via federated learning or directly on the aggregated privacy-protected data. It also allows farmers to identify potential collaborators based on similarities. We have validated this on real-life datasets, demonstrating robust privacy protection against adversarial attacks and utility performance comparable to a centralized system. We demonstrate how this framework can facilitate collaboration among farmers and help researchers pursue broader research objectives. The adoption of the framework can empower researchers and policymakers to leverage agricultural data responsibly, paving the way for transformative advances in data-driven agriculture. By addressing critical privacy challenges, this work supports secure data integration, fostering innovation and sustainability in agricultural systems.

Index Terms—Privacy Enhancing Technologies, Digital Agriculture, Privacy Preserving Framework, Secure Collaborative Research

I. INTRODUCTION

Agriculture has long been the backbone of human civilization, driving innovation to enhance food production and promote sustainability. Over the years, advances such as fertilizers, pesticides, genetically modified seeds, and soil composition analysis have transformed cultivation and harvest practices [1], [2]. Today, this tradition of innovation is driven by a wealth of data collected from satellites, unmanned aerial vehicles, self-driving machines, and field sensors [3]. Although farmers occasionally share sales data for marketing purposes, sharing sensitive data on environmental and soil conditions presents significant privacy challenges [4]. The very data that empower precision agriculture to optimize yields and sustainability also carry risks of misuse, such as adverse pricing, price discrimination, disease interference, inflated insurance costs, and harmful competition [5], [6]. These concerns deter farmers from sharing their data, creating a critical barrier to advancing agricultural research. To unlock the full potential of data-driven agriculture, it is imperative to prioritize protecting farmers from these risks, ensuring that innovation does not come at the expense of their trust and livelihoods.

The existing data collection and collaboration frameworks in the domain, like Farm2Facts [7], [8], Local Line [9], and Open Food Network [10], collect vast agricultural data, but they lack mechanisms for privacy-preserving collaboration. As a result, farmers and researchers are unable to share data freely on such platforms due to proprietary or privacy concerns, which impede innovation. Publicly sharing sensitive agricultural data raises concerns about privacy and security, even for research purposes. Several critical security threats in the digital agriculture domain pose significant risks to farmers' privacy and data security. These threats can be intentional attacks perpetrated by malicious third parties, such as attribute inference attacks and membership inference attacks, or they can be unintentional risks, including configuration errors and improper encryption, which may inadvertently expose sensitive information.

This paper proposes a framework to address these privacy and security challenges in data-driven agriculture. The framework is designed to mitigate the risks of data misuse while

fostering trust among farmers. It achieves this by employing techniques such as differential privacy and dimensionality reduction, which obscure sensitive data and prevent unauthorized reconstruction of the original datasets. Furthermore, the framework aims to reduce the risk of membership inference attacks by ensuring that even anonymized data remains untraceable to individuals. By addressing these specific threats, the framework creates a secure environment for data sharing and collaboration for the farmers, thereby promoting trust and enabling the agricultural research community to fully leverage the potential of data-driven agriculture.

The main goal of this work is the application of privacy-preserving machine learning tools like dimensionality reduction, differential privacy, etc, to address a critical privacy challenge in digital agriculture: enabling collaborative research while safeguarding data privacy. Our research makes a significant contribution by developing a comprehensive privacy-preserving framework for agricultural collaboration. First, we introduce a novel combination of differentially private algorithms (like Laplace noise addition) and dimensionality reduction techniques (like Principal Component Analysis) to safeguard farmers' sensitive data while preserving its utility for collaborative research. Dimensionality reduction algorithms compress the data from a higher dimension to a lower dimension, which obfuscates the original feature space. At the same time, Laplace Noise Addition is a technique used to achieve differential privacy. A combination of these two techniques provides a robust guarantee of security against malicious attacks like membership inference attacks (see details in Section V). Building on this privacy foundation, we then apply clustering algorithms, such as K-Means, to identify and group farmers with similar characteristics, enabling the formation of networks for meaningful collaboration. By grouping farmers based on shared traits, such as crop types, soil composition, or farming practices, the framework facilitates knowledge exchange, the sharing of best practices, and collective problem-solving. This allows farmers to connect with similar farmers based on the similarity of their farm profiles. This creates a collaborative environment that not only fosters peer learning but also empowers farmers to adopt innovative techniques, enhance productivity, and address common challenges through resource sharing and mutual support.

The proposed framework also achieves a robust balance between privacy preservation and utility, enabling effective training of ML models. It ensures that the transformed privacy-protected data aggregated by the sandbox from the farmers preserves the utility required for the effective training of models. Evaluation of the framework on real-world datasets like the Wisconsin Farmer's Market and Crop Recommendation dataset [11] indicates that models trained using the proposed architecture achieve comparable accuracies to those trained on centralized raw data (see details in Section III-D).

Our framework combines privacy preservation with a sandboxed collaboration environment, enabling researchers to derive insights without exposing raw farmer or market data and providing a novel solution for collaborative research in the agricultural domain, where such systems are absent despite widespread demand. We envision it to be deployed as a web-based sandbox environment housed on a cloud computing platform like AWS or Azure, featuring a user-friendly interface and an automated backend for data integration, data processing, model training, and sharing. The design ensures easy adoption by farmers and researchers without requiring technical expertise.

The remainder of this work is structured as follows. To lay the groundwork for our proposed solutions, we start with a literature review of foundational concepts and related works in **Section II: Background and Related Work** . Next, **Section III: Proposed Solution** delves into the design and technical details of our framework. Then, we examine the privacy and utility evaluation metrics and performance of our framework in **Section V: Privacy and Evaluation**. **Section VI: Applications in Machine Learning** explores the application of our framework to real-world scenarios and demonstrates the practical utility of our approach. Finally, **Section VII: Conclusion** summarizes our contributions and discusses their potential impact.

II. BACKGROUND AND RELATED WORK

This section briefly presents literature relevant to our study, including data privacy challenges in digital agriculture, existing frameworks and mechanisms for secure data sharing, and data privacy techniques. Each provides insights into the current state of research and the gaps our work aims to address.

A. Agricultural Privacy Concerns

As the agricultural industry evolves, farmers are increasingly adopting technologies to manage soil, water and harvest more efficiently. Since 2010, these innovations, collectively known as Climate-Smart Agriculture or precision agriculture, have transformed large-scale farming practices by enabling real-time monitoring and input optimization [12], [13]. These practices are critical for commodity farmers, improving resource use efficiency and sustainability by tailoring inputs to specific crop conditions. However, data-driven technologies also risk deepening existing inequalities in the agricultural sector [6], particularly among small and diversified farmers.

Small and medium-sized farmers face significant barriers to adopting data-driven technologies. First, these tools are predominantly designed for large-scale, monoculture farming systems, making them less applicable to diversified farms with varying crops and conditions [14], [15]. Second, while large-scale farmers often share data with industry partners to receive tailored recommendations, small-scale farmers are hesitant due to concerns over potential data misuse [6], [14], [16]. Since 2015, academic research has increasingly recognized these privacy concerns, proposing legislative protections and software solutions to alleviate these challenges [3], [17], [18]. Despite these efforts, as of August 2024, the agricultural sector lacks a dedicated federal regulatory body for data protection—unlike the healthcare (HIPAA) and finance (GBLA) sectors [19]. While the U.S. Cybersecurity and Privacy Operations Center

(CPOC) oversees some agricultural data concerns [20], it lacks the authority to enforce industry-wide protections [19].

Agricultural data storage involves multiple stakeholders, including Agricultural Technology Providers (ATPs), farmers, agronomists, and digital platforms. Emerging technologies, such as edge computing and blockchain, promise robust data security and traceability [21], [22]. However, small-scale farmers face substantial challenges in adopting these advanced systems, often bearing significant risks without receiving proportional benefits [3], [5]. Additionally, the complexity of lengthy, technical data-sharing agreements further discourages smallholder participation, exacerbating their concerns about data misuse and undermining trust in digital agriculture systems [23].

B. Data Sharing Frameworks for Digital Agriculture

The literature on data-sharing frameworks in agriculture has addressed privacy and security concerns. For instance, [24] propose a set of principles and an access control mechanism within Data Sharing Agreements (DSAs) to regulate access to shared data and address privacy concerns in the agricultural sector. Similarly, [25] introduces a flexible, cloud-based, privacy-preserving data aggregation scheme that allows users to access aggregated plain-text data while safeguarding farmers' sensitive information. These frameworks offer scalable solutions for secure data sharing; however, their applicability to small and medium-sized farming operations remains limited.

Most studies focus heavily on securing smart devices from external attacks, often overlooking broader privacy frameworks tailored to diversified farming operations. For example, [26] propose a multi-layered architecture that combines cyber-physical measures to deter cyberattacks on smart farming systems. [27] extend this work by introducing a privacy encoding framework for sensor-based systems, complemented by an intrusion detection system that maintains data integrity. Similarly, [28] developed a privacy-preserving framework for unmanned aerial vehicles (UAVs), emphasizing authentication mechanisms to prevent data poisoning attacks that could compromise agricultural decision-making processes.

Despite these advancements, there is a growing necessity for tailored security frameworks that address the unique challenges faced by small and medium-sized farms, particularly those with diverse farming practices [29]. To the best of our knowledge, few existing frameworks enable effective and privacy-preserving data collection from farmers, particularly for use in AI and machine learning (ML) applications. Our study addresses this critical gap by proposing and testing a privacy-preserving framework that ensures the confidentiality of individual farmers' data while optimizing the utility of aggregated datasets. By balancing privacy and data usability, our framework empowers researchers to develop AI-driven solutions that are both ethical and impactful for agricultural research.

C. Relevant Data Privacy Techniques

In the broader context of data-driven approaches, protecting sensitive information is crucial across various domains, including agriculture. While some data privacy techniques have been widely applied in fields such as finance and healthcare, they also hold potential for enhancing data security in digital agricultural applications. This section examines several relevant key data privacy techniques, exploring both their theoretical foundations and their adaptability for use in agriculture. We will discuss Principal Component Analysis (PCA), which, though not explicitly designed for privacy, offers a form of data abstraction by compressing the data into a lower-dimensional space, thereby obfuscating the original feature space; and Local Differential Privacy (LDP), a method that ensures robust privacy through local data perturbation. Understanding the mathematical principles behind these techniques is essential to appreciating their role in safeguarding sensitive data, particularly as the proposed framework is presented.

Principal Component Analysis (PCA) [30] is a widely utilized technique for data analysis and the identification of patterns and relationships within data. It has a wide variety of applications, such as image compression [31]; facial recognition [32]; medical data correlation [33]; and Quantitative Finance [34]. PCA itself does not inherently provide any guarantees of privacy, as its primary function is to reduce the dimensionality of data by transforming it into a set of orthogonal components. However, this transformation abstracts the data, making it harder to reverse engineer individual records and thereby providing a form of data obfuscation. Similar techniques have been used in the identification of population stratification for genomics research [35].

Local Differential Privacy (LDP) [36], [37] is a privacy-preserving technique that builds upon the principles of classic Differential Privacy (DP) [38]. LDP is designed to protect individual data and extends the concept of DP to a decentralized context. In LDP, each user perturbs their own data prior to sending it to a data collector. This local perturbation ensures that the privacy of individual data is protected, even in cases where the data collector is not considered reliable. Unlike DP, which relies on a central trusted party that has access to the raw data, LDP offers privacy assurances directly at the data source. An algorithm A is considered to satisfy ϵ-LDP if, for any pair of individual private data points a_1 and a_2, and for any possible output b,:

$$\Pr[A(a_1) = b] \leq e^{\epsilon} \cdot \Pr[A(a_2) = b], \qquad (1)$$

where ϵ represents the privacy parameter. ϵ-LDP is achieved by introducing an appropriate amount of noise to individual data points. The primary challenge is determining the noise's magnitude to ensure compliance with LDP while preserving data utility. Various mechanisms have been developed in the DP field to address this problem, with the Laplacian mechanism [39] being one of the most prominent.

For any numerical function $f(x) : \mathbb{R} \rightarrow \mathbb{R}$ that satisfies ϵ-LDP if Laplacian noise is added as follows:

$$F(x) = f(x) + Lap\left(\frac{s}{\epsilon}\right) \qquad (2)$$

where s represents the sensitivity of the function f and Lap(k), where $k = \frac{s}{\epsilon}$, denotes sampling from a Laplace distribution with scale k. Sensitivity measures the maximum change that a single data point can cause to the output of f in the worst-case scenario. In our work, we calculate sensitivity to add Laplacian noise [40]. By definition, $s = \max_{x,x'} \|f(x) - f(x')\|_1$.

III. PROPOSED SOLUTION

A. System Model

Our proposed privacy-preserving comprises two main actors: researchers R and participants P_i (e.g., farmers). Participants generate, own, and retain their data to derive insights and benefits from it. Researchers, on the other hand, need the data to conduct research analysis and machine learning model training (see Table I). The relationship between researchers and participants is inherently symbiotic as researchers need data to do their work and ultimately, their work provides benefits to the farmers in the form of machine learning models such as predictive models for disease detection or weather forecasting. The framework is a sandbox environment that performs key operations, such as data aggregation, clustering, and filtering in the middle. Together, these components form the complete the entire architecture.

B. Threat Model

In the proposed framework we assume the farmers to be trusted entities as they provide data to the researcher. They generate private agricultural data and perform local PCA transformations and noise addition (to achieve differential privacy for the shared data) before sharing the transformed data.

We assume the researcher to be an honest but curious party. They follow the protocol's specifications correctly. However, they might be curious to learn unauthorized information about farmers' private data through protocol outputs, which include analyzing the transformed data matrices, clustering results, and machine learning models. The researcher has access to public agricultural datasets and can perform polynomial-time computations on all protocol outputs.

We also assume the server to be honest but curious. Nonetheless, the protocol protects against data leakage to the server by implementing differential privacy techniques. These techniques add noise to the data, obscuring individual farmers' sensitive information. The protocol also utilizes dimensionality reduction methods, such as PCA, to further obfuscate the data structure, making it challenging for attackers to extract meaningful insights from the interactions they view.

We consider various adversarial attacks like membership inference attacks, attribute inference attacks, and deanonymization attacks. Since the identities of the participants who share the data are hidden, attribute inference attacks are only feasible if an attacker can infer the victim's membership to the research

Notation	Description
R	Researcher
P_i	Participant i
D_s	Dataset used by the researcher to train the PCA model
M_s	Trained PCA model
O_i	Original PCA matrix of participant P_i
C_i	Noisy PCA matrix of participant P_i

dataset through a membership inference attack. The proposed protocol is designed to be robust against membership inference attacks through several mechanisms. We evaluate the robustness of our framework against inference attacks by performing power analysis (see Section V). Although deanonymization attacks can be effective when auxiliary data is available, obscuring original feature space using dimensionality reduction and applying noise to shared datasets significantly reduces its success.

Each farmer independently performs data transformation and noise addition, ensuring that compromising one farmer's process cannot affect others. The framework infrastructure enforces a strict separation between different farmers' data streams, preventing unauthorized mixing or correlation of data from various sources. The framework infrastructure is assumed to maintain its integrity and not collude with any party to breach privacy guarantees.

C. Proposed Scheme

We propose a privacy-preserving mechanism using dimensionality reduction and differential privacy techniques to identify farmers having a specific set of attributes. Figures 1 and 2 illustrate the framework design and workflow. Figure 1 illustrates a high-level conceptual overview of the framework's architecture and workflow, free from technical distractions. It provides a high-level system figure, and it emphasizes the end-to-end process, focusing on roles, data flow, and anonymization at an abstract level to enhance clarity for the audience. Figure 2 presents a more detailed overview of the proposed framework, concentrating on the granular technical aspects of the core components, data flow, data processing, and infrastructure details essential for replicability and a deeper understanding. See Table I for the list of notations and their definitions.

The researcher R trains a global dimensionality reduction model such as a principal component analysis (PCA) model on a publicly available dataset D_s having desired farm features (e.g., nitrogen levels, annual rainfall, or pH). The global model is generated to ensure that participants can map their locally reduced data into a common space, allowing for consistent and comparable results when the data is integrated or analyzed together. We use the Standard Scaler function from scikit-learn [41] to preprocess dataset D_s. This trained dimensionality reduction model M_s is then shared with all the participants through the sandbox F along with the sensitivity parameter, which will be used later in the noise addition step. Each P_i

uses model M_s to transform its private data into transformed matrix O_i. To guarantee the privacy of the data, especially against membership inference attacks, the framework applies privacy enhancement techniques like adding Laplacian noise locally to achieve ϵ-LDP. Noise is added to each component of a O_i matrix using 1, where sensitivity s is replaced with s_i, which denotes the sensitivity of component i in the O_i matrix. Finally, a differentially private version of their transformed data C_i is sent to sandbox F.

The sandbox environment receives differentially private data collected from multiple farmers, aggregates it, and then applies clustering algorithms to group farmers. Figure 2 presents the complete workflow of the framework, with a specific use case of connecting similar farmers based on farm attributes. Farmers interact with the sandbox by uploading their farm profile, which is then input to a clustering algorithm, such as K-Means [42], pre-trained on aggregated data. The clustering labels are utilized to group farmers and recommend connections among those with similar characteristics. Any P_i (farmer) can query the sandbox with their farm profile to identify potential collaborators using aggregated data. Farmers can communicate, share knowledge, and collaborate with recommended potential collaborators.

Similarly, the researcher R can utilize the aggregated data to query the sandbox with a target farm profile (e.g., similar to a farm P_i) to identify potential collaborators. This helps the researcher R to train various machine learning (ML) models on identified farmers' data. R identifies potential collaborators of a target P_i and uses Federated Learning (FL) techniques to train personalized machine learning models tailored to target P_i profile trained on data from identified potential collaborators. Furthermore, machine learning models can be trained directly on the aggregated data to make predictions. These models can predict various outcomes, such as the likelihood of farm success or failure, or analyze risks like disease outbreaks. It is important to emphasize that R does not have direct access to the private datasets. All analyses and predictions are conducted exclusively using differentially private outputs collected from farmers, ensuring a completely privacy-preserving approach.

D. Case Study

For feasibility testing and evaluation[1], the following case study explores the potential of the proposed privacy-preserving framework for connecting farmers based on attributes and evaluates it using two distinct datasets: a Wisconsin Farmer's Market dataset (Table II) and a crop recommendation dataset [11] (Table III). We consider the following use case from the domain of digital agriculture research: The farmers seek to identify potential collaborators based on farm characteristics to collaborate.

Farmer's Market dataset (Table II), which was gathered internally and has not been made public, includes detailed

---

[1]All implementation codes are available on GitHub and will be provided upon request.

information about various aspects of farmer's market activities and vendors and was explicitly collected for research purposes. Due to the private nature of the information contained in this dataset, it remains private. Meanwhile, the crop recommendation dataset (Table III) is a publicly available dataset collected from multiple farms representing the soil composition and environmental factors [11]. The dataset is divided into five pieces representing various markets with uneven sample distributions for each crop, collecting data from different farmers and sharing it with researchers. We consider a researcher receiving the dataset in a privacy-preserving manner using the aforementioned techniques and applying a clustering algorithm like K-Means. Figure 3 and Figure 4 illustrate the plot of K-Means clustering on the PCA-transformed farmer's market dataset in Table II and crop recommendation dataset in Table III, respectively. While recommending potential collaborators, the farm profile is fed to the K-Means model to get the group label for the farm. Then, we use the nearest neighbor algorithm [43] to find the nearest neighbors of the farm profile having the same group label.

The entire process is repeated at the Farmer's Market dataset level to identify potential collaborators for personalized machine learning model training via federated learning (see Figure 1, which illustrates the end-to-end lifecycle of the framework, showcasing one of its use cases: how the researcher uses aggregated data to identify partners and train personalized models.). The crop recommendation dataset (Table III) is divided into a global training dataset and multiple farmer's market datasets to simulate experiments. Figure 5 illustrates the plot of K-Means clustering on the PCA-transformed crop recommendation dataset in Table III divided into five distributed farmer's markets and a global dataset. For personalized model training, a participant P_i (framer's markets in this case) initiates the model training. Frameworks identify three potential collaborators of the initiator P_i out of 5 potential collaborators who can contribute to and participate in training machine learning models under a federated learning setting, enabling secure and collaborative research. There are two ways for a framework to identify potential collaborators: (i) based on the similarity of the participant's profile (data) with each other, and (ii) based on the similarity of distribution of participant's data. Framework uses the nearest neighbor algorithm [43] to sort all the participants in descending order of their similarity and selects the first three. The trained model is personalized for the initiator participant. Figure 6 and Figure 7 illustrate the accuracy of the federated learning model vs. the epsilon privacy budget for a neural network trained on three identified collaborators of the initiator P_i. In Figure 6, the collaborators are selected based on the similarity of their participant profiles, while in Figure 7, the collaborators are chosen based on the similarity of the distribution of their data.

IV. SECURITY ANALYSES

The security of the proposed framework reduces to the security of the underlying ϵ-LDP mechanism. In the following,
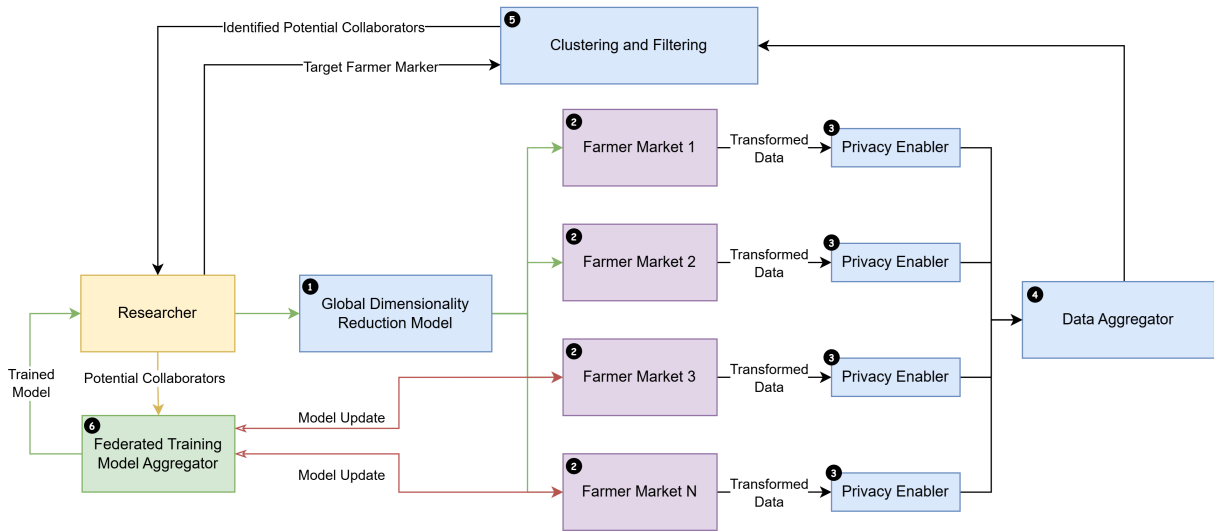
Fig. 1. A high-level data flow: the researcher sends a globally trained dimensionality reduction model to N farmer markets, who transform and anonymize their data before sending it to a data aggregator for sharing on the sandbox platform. The researcher then identifies potential collaborators to train a personalized model.
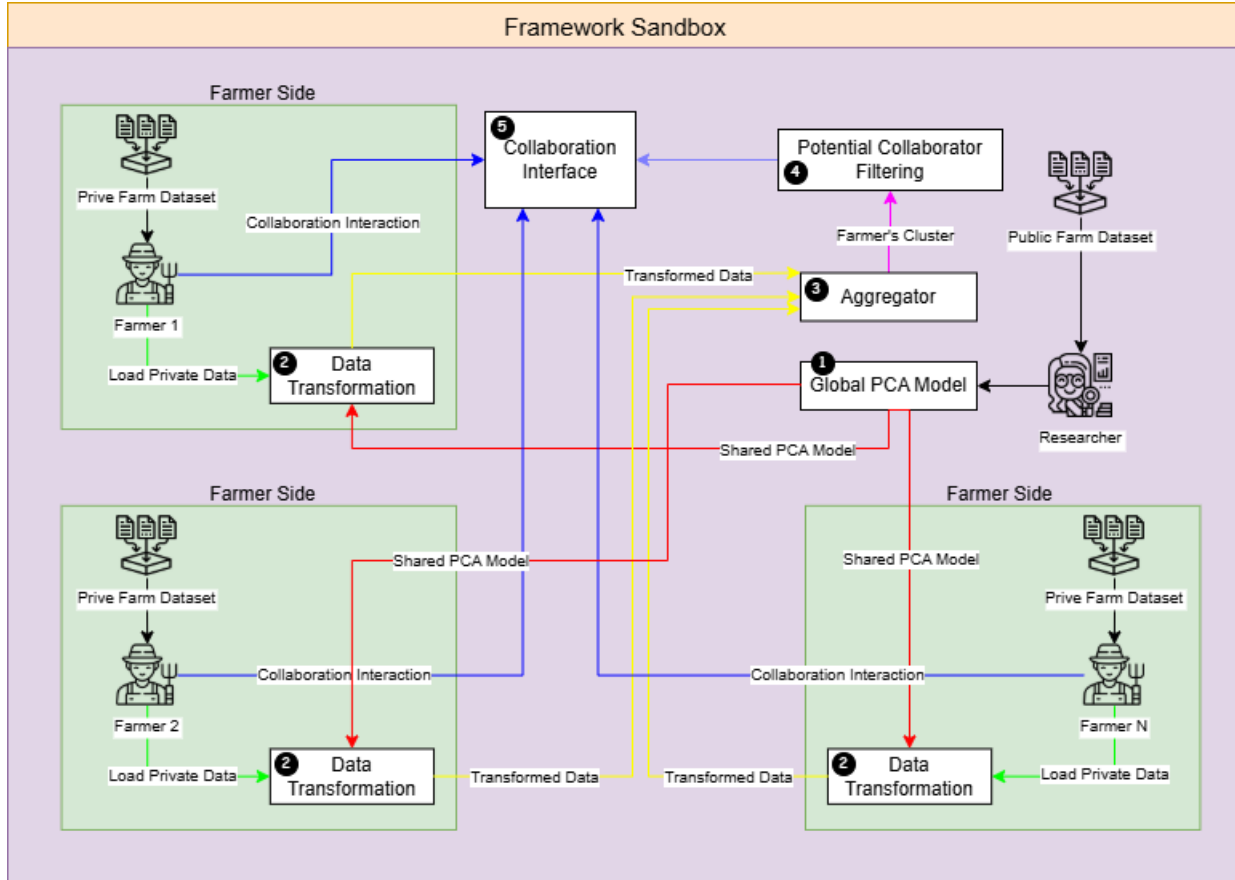

Fig. 2. Detailed architecture of the proposed framework.

we provide the formal security theorem and its proof.

Theorem 1. *Let M_s be the randomization mechanism (e.g., Laplacian mechanism) used in the reduction model. The*

privacy-preserving scheme for the collaborative agriculture research sharing framework is secure and protects the privacy of the farmers against PPT semi-honest adversaries if the

TABLE II
FARMERS MARKET DATASET.

	Miles from Market	Vendor Type									Sales	#Visitors
ID	Primary Location	F&V	M&S	D	Eggs	P&F	N&L	VA	PF	CAS	Sales	#Visitors
6815	42.22	0	0	1	0	0	0	1	1	0	414.0	70
5991	19.62	1	1	0	0	0	0	0	0	0	157.0	404
5663	23.54	0	0	0	0	0	0	0	0	1	670.0	6
5950	70.52	0	0	0	0	0	0	1	0	0	70.0	53
5686	0.68	0	0	0	0	0	0	0	0	0	285.0	166

F&V: Fruits and Vegetables; M&S Meat and Seafood; D:Dairy; P&F: Plants and Flowers; N&L: Nuts and Legumes; VA: Value-added; PF: Prepared Food; CAS: Crafts/Art/Services.

TABLE III
CROP RECOMMENDATION DATASET.

N	P	K	Temp	H	pH	R	Label
90	42	43	20.879	82.002	6.502	202.935	rice
61	44	17	26.100	71.574	6.931	102.266	maize
40	72	77	17.024	16.988	7.485	88.551	chickpea
13	60	25	17.1636	20.595	5.685	128.256	kidney beans
90	46	42	23.978	81.450	7.502	250.082	rice

N: Nitrogen; P: Phosphorous; K:Potassium; pH: pH value; R: rainfall in mm; Temp: temperature in Celsius; H: relative humidity in %.
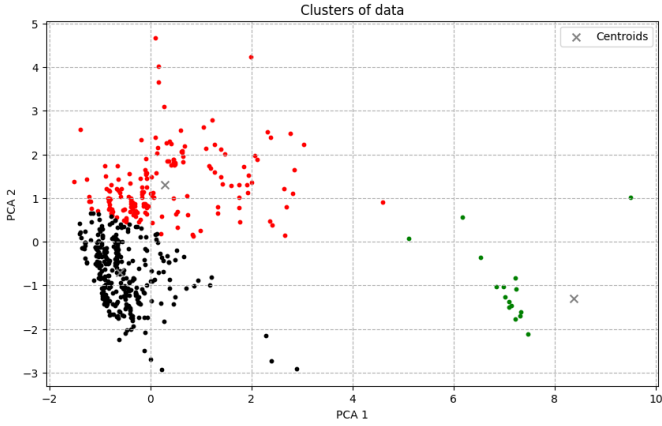
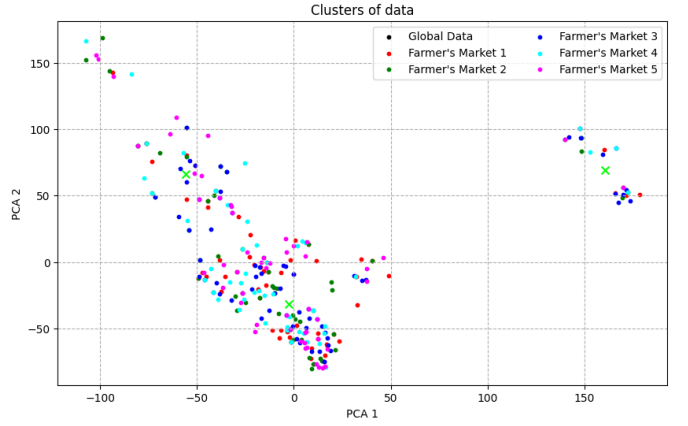Fig. 3. PCA transformed farmer's market dataset clustering.

Fig. 5. K-Means clustering on the PCA-transformed crop recommendation dataset, grouped into five farmer's markets and a global set.

Fig. 4. PCA transformed crop recommendation dataset clustering.

PCA transformation and Laplacian noise satisfies ϵ-Local Differential Privacy.

We prove theorem 1 by introducing two worlds. There is an ideal world the adversary's view $\text{View}_A^{\text{Ideal}}$ is the outputs of the interactions with a simulator who has no access to the datasets. Moreover, in the real world, the adversary's view $\text{View}_A^{\text{Real}}$ is the actual output of the framework using the reduction model and Laplacian noise mechanism. The ϵ-LDP guarantee ensures that the outputs of the mechanism for neighboring datasets are statistically close (up to a factor of e^ϵ), meaning that the adversary's view in the ideal world is computationally indistinguishable from the real world view.

$$\text{View}_A^{\text{Real}} \overset{c}{\approx} \text{View}_A^{\text{Ideal}}$$
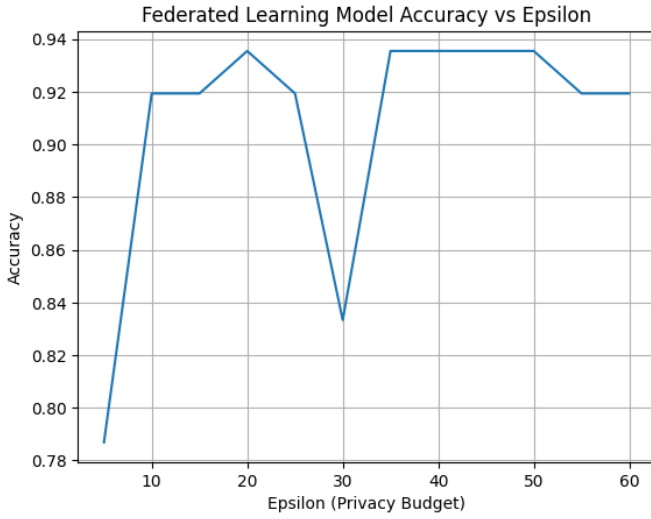
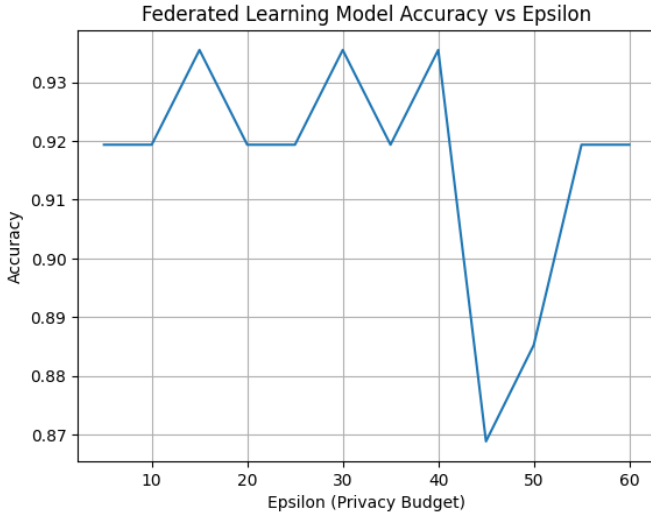Fig. 6. Federated learning accuracy vs. epsilon (based on profile similarity).

Fig. 7. Federated learning accuracy vs. epsilon (based on data distribution similarity)

Where $\stackrel{c}{\approx}$ indicates that no PPT adversary A can distinguish between the two views with non-negligible advantage.

Proof. In the ideal world, the simulator S accesses the privacy parameter ϵ and the adversary's queries and cannot access the actual datasets D_s^0 or D_s^1. For each query made by the adversary, S generates noisy outputs using an ideal ϵ-LDP mechanism. The outputs are indistinguishable from those in the real world because of the ϵ-LDP guarantee.

The real world interacts with the adversary A by answering their queries and providing noisy outputs according to the actual interactions of the privacy-preserving framework.

If A has a non-negligible advantage in distinguishing between the two worlds, we can construct a distinguisher D that breaks the ϵ-LDP guarantee of the mechanism M_s.

The distinguisher D receives the output C_i generated from M_s for one of two neighboring datasets D_s^0 or D_s^1 and simulates the real-world mechanism using C_i as the output of M_s. Then, it interacts with the adversary A to send A's queries to M_s, generating noisy outputs as needed. D observes A's decision and uses it to distinguish which dataset C_i was generated from. If A can distinguish between the real and ideal worlds with a non-negligible advantage, then D can distinguish between the outputs of M_s applied to D_s^0 and D_s^1, contradicting the ϵ-LDP guarantee.

By contradiction, A cannot distinguish between the real and ideal worlds with non-negligible advantages. Therefore, the outputs of the real-world mechanism are computationally indistinguishable from those of the ideal-world mechanism, and the privacy-preserving framework is secure under the ϵ-LDP assumption, which concludes the proof. □

V. PRIVACY AND UTILITY EVALUATION

The farmer data aggregated in a privacy-preserving manner using the proposed framework can be used to train clustering and the nearest neighbor models to find a particular farmer's nearest neighbors (potential collaborators). A researcher can identify potential collaborators and train an ML model on their raw data in a privacy-preserving manner using techniques like Federated Learning. To evaluate the performance of a model trained on data of identified potential collaborators of a particular farmer, we compare its performance with the performance of a model trained on all farmer datasets.

We consider the Wisconsin Farmer's Market dataset (See Table II) from our case study section. To simulate two local farmer's market datasets P_is and the dataset used by the researcher to train the PCA model D_s, we divide our dataset into three parts. Each P_i converts their data to C_i using shared PCA model M_s, which is accessible to the sandbox.

To evaluate the utility of our process, we compare the accuracies of a classifier trained on raw data aggregated from all the farmers in a centralized manner vs a classifier model trained on data aggregated through the proposed architecture. Table IV illustrates the accuracies of various models trained under both settings. The models trained using the proposed architecture exhibit performance comparable to that of the traditionally trained models on centralized data, with an average accuracy loss of 5.33%. This difference represents a reasonable trade-off, given the substantial privacy benefits achieved.

To evaluate the effect of noise addition on the utility of the data, we compared model performance on aggregated data with and without noise across varying values of epsilon (ϵ). The parameter ϵ quantifies privacy protection, where smaller values of ϵ correspond to stronger privacy guarantees and higher noise, leading to a trade-off with utility. Figures 8, 9, and 10 illustrate the power (fraction of correctly identified noisy samples) and accuracy for a range of ϵ values for the Logistic Regression, Naive Bayes, and Support Vector Machine (SVM) algorithms, respectively on crop recommendation dataset.

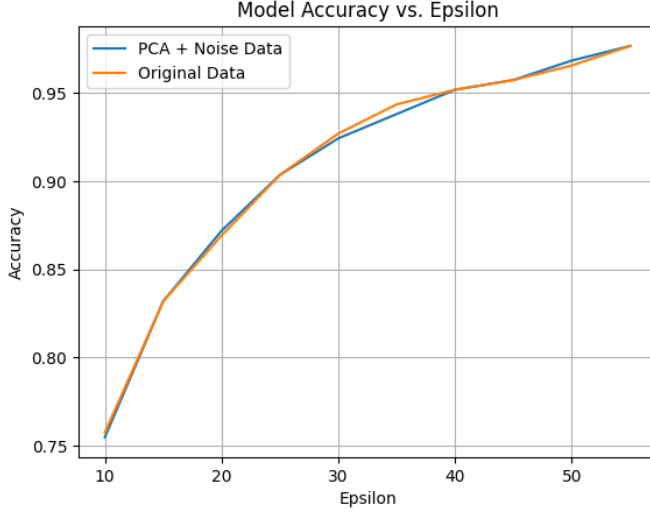| Model Name	Accuracy on Centralized Data	Accuracy on Aggregated Data
Logistic Regression	99.6	92.1
Naive Bayes	98.3	93.8
Support Vector Machine	99.5	95.5

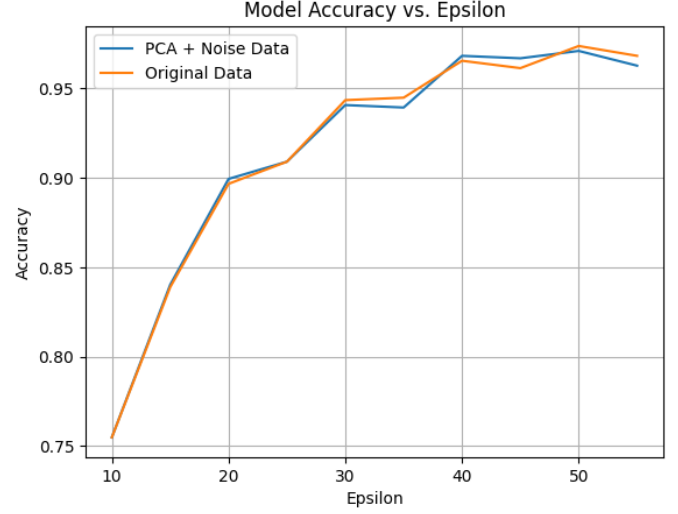Fig. 8. Model accuracy on noisy and original data (Logistic Regression) vs epsilon on crop recommendation dataset.

Fig. 10. Model accuracy on noisy and original data (SVM) vs epsilon on crop recommendation dataset.
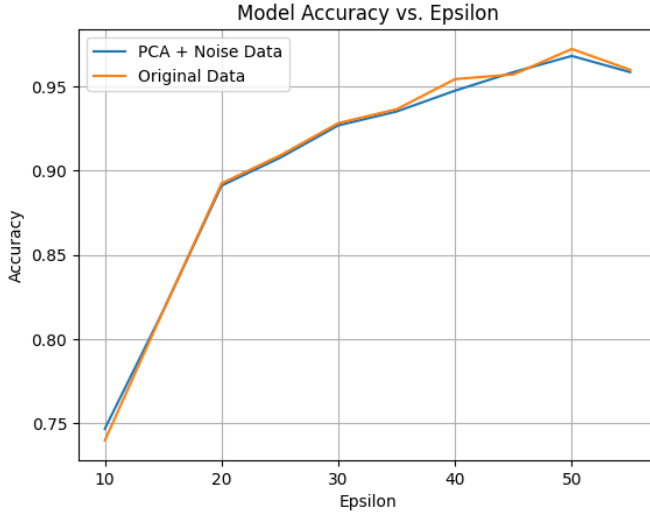

Fig. 9. Model accuracy on noisy and original data (Naive Bayes) vs epsilon on crop recommendation dataset.

To evaluate the privacy performance of our algorithm, we compute the power metric, which indicates the fraction of correctly identified noisy samples. To achieve this, we calculate the minimum pairwise distances between the samples of the shared dataset and the control group as the control distances. Similarly, we compute the minimum pairwise distance of the shared dataset and case group as case distances. Here, the control group is sampled from the sandbox global dataset, and the case group is sampled from the participant's shared dataset. Then, we determine the threshold, defined as the distance at the 95th percentile (when the False Positive Rate is 5%) of the sorted control distances. This threshold is subsequently used to evaluate the fraction of case distances that fall below it, thereby determining the power. The metric is evaluated with specific thresholds for False Positive Rate (FPR) and epsilon (ϵ). ϵ is a privacy parameter of privacy loss at a differential alteration in the dataset and is allocated across columns according to their sensitivity. For our analysis, we set the FPR at 0.05 (5%) and incrementally vary the ϵ value, starting from 10, to observe its impact on the power metric. This approach helps us select the correct value of the ϵ with acceptable robustness represented by the power metric.

For the utility metric, we compare the noise-free cluster labels of the data with the labels predicted by a binary classifier for noisy samples, assessing the accuracy of correctly predicting the labels of noisy samples. We have performed the aforementioned utility test using three different algorithms (Logistic Regression, Naive Bayes, and Support Vector Machine) to demonstrate the results across classifiers. Figures 14, 15, and 16 illustrate the power and accuracy for a range of ϵ values for the Logistic Regression, Naive Bayes, and SVM algorithms, respectively, on the farmer's market dataset. Similarly, Figures 11, 12, and 13 illustrate the power and
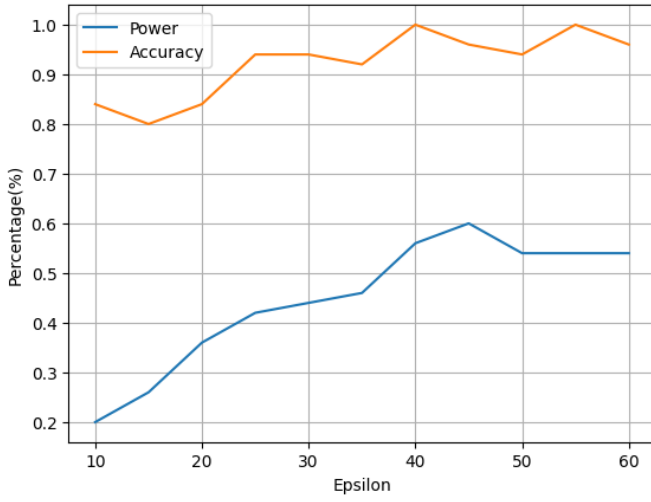
Fig. 11. Power and accuracy (Logistic Regression) vs epsilon on crop recommendation dataset.
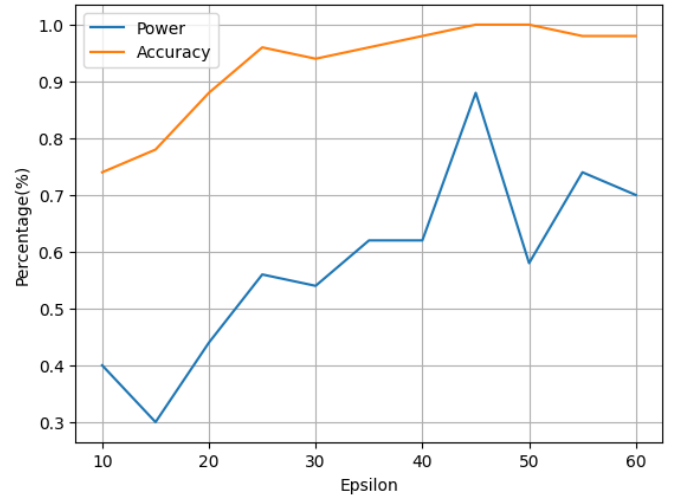

Fig. 13. Power and accuracy (SVM) vs epsilon on crop recommendation dataset.
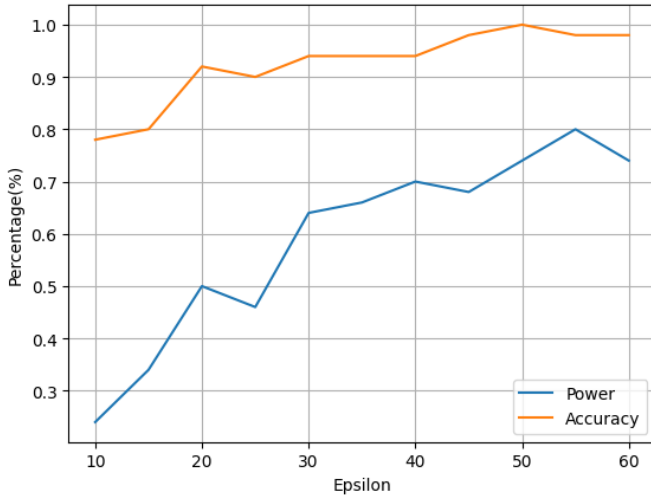

Fig. 12. Power and accuracy (Naive Bayes) vs epsilon on crop recommendation dataset.
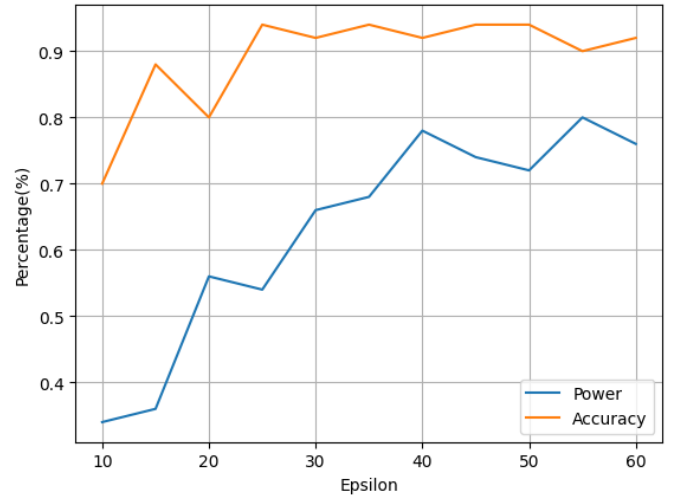

Fig. 14. Power and accuracy (Logistic Regression) vs epsilon on farmer's market dataset.

accuracy for a range of ϵ values for the Logistic Regression, Naive Bayes, and SVM (Support Vector Machine) algorithms, respectively on crop recommendation dataset. The results reflect the high robustness of privacy with comparable utility accuracy. For each model, the optimal ϵ value is determined based on a trade-off between privacy (measured by the power metric) and utility (measured by accuracy). Specifically, the optimal ϵ values correspond to those that achieve both low power (indicating strong privacy protection) and high accuracy. Therefore, optimal epsilon values are 25, 35, and 35 for the Logistic Regression, Naive Bayes, and Support Vector Machine models, respectively.

Overall, the framework achieves comparable model accuracies to centralized data training, with minimal utility loss across diverse classifiers (see Table IV). This demonstrates that privacy-preserving techniques can be effectively integrated into machine learning models without compromising performance. Operating under realistic assumptions of non-colluding participants and honest-but-curious researchers, a common model in privacy-preserving machine learning, the framework significantly mitigates the effectiveness of adversarial attacks like membership inference attacks, providing a robust defense against privacy breaches. The use of the PCA algorithm transforms data into principle components, thereby obfuscating the original feature space and protecting the original data from reconstruction attacks.

Additionally, the analysis of noise addition and the power metric further confirm that optimal ϵ values yield comparable accuracies across various machine learning models, demonstrating the framework's robustness against adversarial attacks and underscores for real-world collaborative machine learning applications.
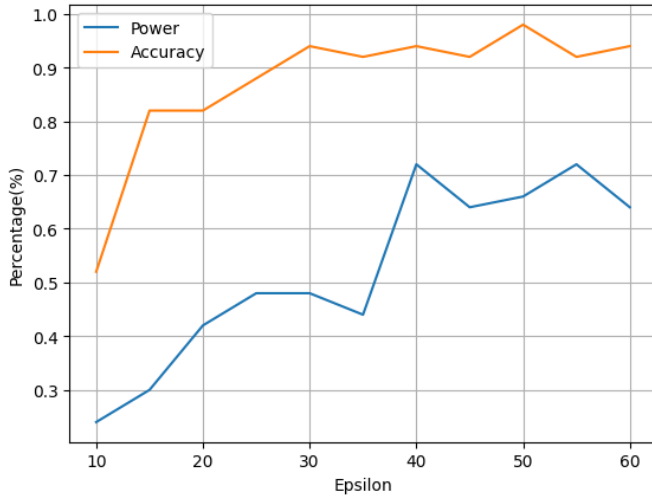
Fig. 15. Power and accuracy (Naive Bayes) vs epsilon on farmer's market dataset.
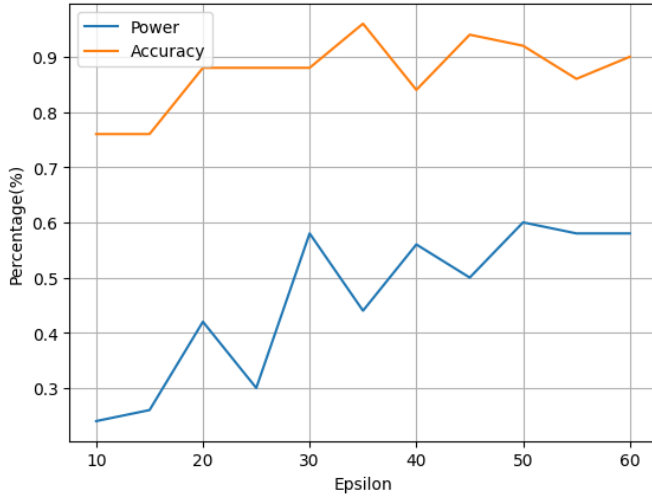

Fig. 16. Power and accuracy (SVM) vs epsilon on farmer's market dataset.

VI. Applications in Machine learning

The proposed framework enables researchers to aggregate private data from multiple farmers to train a diverse array of machine-learning models, thereby expanding the scope and quality of their research. We explore a few examples based on the case study dataset that researchers can use to train machine learning models: Researchers can train a model to predict the ideal crop type for a farm based on factors such as soil composition, humidity, and temperature from the market. Additionally, they can utilize farmer's market datasets to develop a model for predicting farm success based on size, distance to market visitors, sales, etc. Alternatively, they can train a machine learning model to predict food accessibility based on factors such as distance from the market, visitor logs, and general public insecurity levels. Similarly, modeling and predicting soil condition data of farms over time to propose management practices and predict crop yields in the region; collecting weather data with soil quality, soil moisture, and crop water requirements to train models for optimizing irrigation systems and predicting water needs in the region; and assessing the impact of climate change on crop yield and soil quality.

VII. Conclusion

Privacy concerns remain a significant barrier to sharing sensitive data, particularly when such data could be misused against its owner. This is especially true for farmers, who are often reluctant to share farm data despite its critical role in driving innovation and research in precision agriculture. A robust, privacy-preserving data-sharing solution is, therefore, essential to balance the need for innovation with the protection of sensitive information. This study presents a novel framework that effectively addresses this challenge by employing techniques such as Principal Component Analysis (PCA) and differential privacy. The framework safeguards individual privacy while enabling researchers to leverage aggregated data for machine learning (ML) and AI applications without compromising data utility.

Our contributions advance the field by introducing a privacy-preserving framework that combines dimensionality reduction with differential privacy techniques to enhance both security and data usability. Additionally, the framework facilitates the identification and grouping of farmers with shared characteristics through clustering algorithms, fostering collaboration and collective problem-solving among farmers with similar needs. By addressing these critical gaps, this work provides a scalable and ethical solution for advancing AI-driven research while empowering the agricultural community.

Despite its strengths, the framework has limitations. For instance, the computational requirements of differential privacy techniques may pose challenges for widespread adoption among resource-constrained farming operations. Additionally, the framework's reliance on consistent and high-quality input data may limit its effectiveness in regions where data collection is inconsistent or less accessible. Future work will focus on refining the framework to enhance its scalability and adaptability to diverse agricultural systems.

In addition, further research will explore integrating decentralized data-sharing mechanisms, such as blockchain, and incorporating real-time data processing to improve trust and responsiveness. By addressing these challenges, this work represents an important step toward enabling privacy-preserving and equitable digital agriculture, laying a foundation for sustainable and ethical farming practices that benefit both researchers and farmers.

VIII. Acknowledgments

REFERENCES

[1] J. Pretty, "Agricultural sustainability: concepts, principles and evidence," *Philosophical Transactions of the Royal Society B: Biological Sciences*, vol. 363, no. 1491, pp. 447–465, 2008.

[2] D. Tilman, K. G. Cassman, P. A. Matson, R. Naylor, and S. Polasky, "Agricultural sustainability and intensive production practices," *Nature*, vol. 418, no. 6898, pp. 671–677, 2002.

[3] S. Wolfert, L. Ge, C. Verdouw, and M.-J. Bogaardt, "Big data in smart farming–a review," *Agricultural systems*, vol. 153, pp. 69–80, 2017.

[4] C. DeMets and A. Morales, "Marketplace measurement: farmers, farmers markets and ecosystem services," *International Journal of Sociology and Social Policy*, vol. 42, 03 2022.

[5] L. Wiseman, J. Sanderson, A. Zhang, and E. Jakku, "Farmers and their data: An examination of farmers' reluctance to share their data through the lens of the laws impacting smart farming," *NJAS-Wageningen Journal of Life Sciences*, vol. 90, p. 100301, 2019.

[6] M. Taylor, "Climate-smart agriculture: what is it good for?" *The Journal of Peasant Studies*, vol. 45, no. 1, pp. 89–107, 2018.

[7] "Farm2facts." [Online]. Available: https://farm2facts.org/

[8] E. Ledesma, A. He, P. Warsaw, L. Suerth, A. Morales, L. Rosenblum, and B. Wiedenfeld, "Citizen scientist: Farm 2 facts supporting farmers markets," *Sustainability*, vol. 13, no. 11, 2021. [Online]. Available: https://www.mdpi.com/2071-1050/13/11/6162

[9] "Local line: The farm to fork commerce platform." [Online]. Available: https://www.localline.co/

[10] "The open food network." [Online]. Available: https://openfoodnetwork.org/

[11] "Crop recommendation dataset," https://www.kaggle.com/datasets/atharvaingle/crop-recommendation-dataset, accessed: 2024.

[12] FAO, "Fao.org," 2024. [Online]. Available: https://www.fao.org/climate-smart-agriculture/overview/faqs/history/en/

[13] J. Lowenberg-DeBoer and B. Erickson, "Setting the record straight on precision agriculture adoption," *Agronomy Journal*, vol. 111, no. 4, pp. 1552–1569, 2019.

[14] R. B. Wakweya, "Challenges and prospects of adopting climate-smart agricultural practices and technologies: Implications for food security," *Journal of Agriculture and Food Research*, vol. 14, p. 100698, 2023. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S2666154323002053

[15] L. Klerkx, E. Jakku, and P. Labarthe, "A review of social science on digital agriculture, smart farming and agriculture 4.0: New contributions and a future research agenda," *NJAS-Wageningen journal of life sciences*, vol. 90, p. 100315, 2019.

[16] O. Westermann, W. Förch, P. Thornton, J. Körner, L. Cramer, and B. Campbell, "Scaling up agricultural interventions: Case studies of climate-smart agriculture," *Agricultural Systems*, vol. 165, pp. 283–293, 2018. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0308521X1830043X

[17] S. Linsner, F. Kuntke, E. Steinbrink, J. Franken, and C. Reuter, "The role of privacy in digitalization–analyzing perspectives of german farmers," *Proceedings on Privacy Enhancing Technologies*, 2021.

[18] J. Kaur, S. M. Hazrati Fard, M. Amiri-Zarandi, and R. Dara, "Protecting farmers' data privacy and confidentiality: Recommendations and considerations," *Frontiers in Sustainable Food Systems*, vol. 6, 2022. [Online]. Available: https://www.frontiersin.org/journals/sustainable-food-systems/articles/10.3389/fsufs.2022.903230

[19] J. L. Ferris, "Data privacy and protection in the agriculture industry: Is federal regulation necessary?" *Minn. JL Sci. & Tech.*, vol. 18, p. 309, 2017.

[20] 2024. [Online]. Available: https://www.usda.gov/ocio/centers/cpoc

[21] M. Amiri-Zarandi, R. A. Dara, E. Duncan, and E. D. Fraser, "Big data privacy in smart farming: a review," *Sustainability*, vol. 14, no. 15, p. 9120, 2022.

[22] A. Kamilaris, A. Fonts, and F. X. Prenafeta-Boldύ, "The rise of blockchain technology in agriculture and food supply chains," *Trends in food science & technology*, vol. 91, pp. 640–652, 2019.

[23] K. Kosior, "Digital transformation in the agri-food sector–opportunities and challenges," *Roczniki (Annals)*, vol. 2018, no. 2, 2018.

[24] K. Spanaki, E. Karafili, and S. Despoudi, "Ai applications of data sharing in agriculture 4.0: A framework for role-based data access control," *International Journal of Information Management*, vol. 59, p. 102350, 2021. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0268401221000438

[25] J. Song, Q. Zhong, W. Wang, C. Su, Z. Tan, and Y. Liu, "Fpdp: Flexible privacy-preserving data publishing scheme for smart agriculture," *IEEE Sensors Journal*, vol. 21, no. 16, pp. 17430–17438, 2021.

[26] M. Gupta, M. Abdelsalam, S. Khorsandroo, and S. Mittal, "Security and privacy in smart farming: Challenges and opportunities," *IEEE Access*, vol. 8, pp. 34564–34584, 2020.

[27] P. Kumar, G. P. Gupta, and R. Tripathi, "Pefl: Deep privacy-encoding-based federated learning framework for smart agriculture," *IEEE Micro*, vol. 42, no. 1, pp. 33–40, 2022.

[28] R. Kumar, P. Kumar, R. Tripathi, G. P. Gupta, T. R. Gadekallu, and G. Srivastava, "Sp2f: A secured privacy-preserving framework for smart agricultural unmanned aerial vehicles," *Computer Networks*, vol. 187, p. 107819, 2021. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1389128621000086

[29] P. A. Ongadi, "A comprehensive examination of security and privacy in precision agriculture technologies," *GSC Advanced Research and Reviews*, vol. 18, no. 1, pp. 336–363, 2024.

[30] K. P. F.R.S., "Liii. on lines and planes of closest fit to systems of points in space," *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, vol. 2, no. 11, pp. 559–572, 1901.

[31] V. S. V. Gaidhane and M. Kumar, "Image compression using pca and improved technique with mlp neural network," *IEEE International Conference on Advances in Recent Technologies in Communication and Computing*, pp. 106–110, 2010.

[32] R. Gottumukkal and V. Asari, "An improved face recognition technique based on modular pca approach," *Pattern Recognition Letters*, vol. 25, pp. 429–436, 03 2004.

[33] N. Qureshi, V. Suthar, H. Magsi, M. Sheikh, M. Pathan, and B. Qureshi, "Application of principal component analysis (pca) to medical data," *Indian Journal of Science and Technology*, vol. 10, pp. 1–9, 02 2017.

[34] H. Yu, R. Chen, and G. Zhang, "A svm stock selection model within pca," *Procedia Computer Science*, vol. 31, p. 406–412, 12 2014.

[35] L. Dervishi, W. Li, A. Halimi, X. Jiang, J. Vaidya, and E. Ayday, "Privacy preserving identification of population stratification for collaborative genomic research," *Bioinformatics*, vol. 39, no. Supplement_1, pp. i168–i176, 06 2023. [Online]. Available: https://doi.org/10.1093/bioinformatics/btad274

[36] J. Duchi, M. Jordan, and M. Wainwright, "Local privacy and statistical minimax rates," *Foundations of Computer Science, 1975., 16th Annual Symposium on*, 02 2013.

[37] P. Kairouz, S. Oh, and P. Viswanath, "Extremal mechanisms for local differential privacy," vol. 17, 04 2016.

[38] C. Dwork, "Differential privacy: A survey of results," vol. 4978, 04 2008, pp. 1–19.

[39] C. Dwork, F. McSherry, K. Nissim, and A. Smith, "Calibrating noise to sensitivity in private data analysis," vol. Vol. 3876, 01 2006, pp. 265–284.

[40] C. Dwork and A. Roth, "The algorithmic foundations of differential privacy," *Foundations and Trends in Theoretical Computer Science*, vol. 9, 01 2013.

[41] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and Édouard Duchesnay, "Scikit-learn: Machine learning in python," *Journal of Machine Learning Research*, vol. 12, no. 85, pp. 2825–2830, 2011. [Online]. Available: http://jmlr.org/papers/v12/pedregosa11a.html

[42] J. A. Hartigan and M. A. Wong, "Algorithm as 136: A k-means clustering algorithm," *Journal of the Royal Statistical Society. Series C (Applied Statistics)*, vol. 28, no. 1, pp. 100–108, 1979. [Online]. Available: http://www.jstor.org/stable/2346830

[43] T. Cover and P. Hart, "Nearest neighbor pattern classification," *IEEE Transactions on Information Theory*, vol. 13, no. 1, pp. 21–27, 1967.