



Published in final edited form as:

Proc IEEE Symp Secur Priv. 2023 May ; 2023: 1908–1925. doi:10.1109/sp46215.2023.10179350.

Scalable and Privacy-Preserving Federated Principal Component Analysis

David Froelicher^{*,†,‡}, Hyunghoon Cho^{*,‡}, Manaswitha Edupalli[‡], Joao Sa Sousa[§], Jean-Philippe Bossuat[¶], Apostolos Pyrgelis[§], Juan R. Troncoso-Pastoriza[¶], Bonnie Berger[†], Jean-Pierre Hubaux^{§,¶}

[†]MIT

[‡]Broad Institute of MIT and Harvard

[§]EPFL

[¶]Tune Insight SA

Abstract

Principal component analysis (PCA) is an essential algorithm for dimensionality reduction in many data science domains. We address the problem of performing a federated PCA on private data distributed among multiple data providers while ensuring data confidentiality. Our solution, SF-PCA, is an end-to-end secure system that preserves the confidentiality of both the original data and all intermediate results in a passive-adversary model with up to all-but-one colluding parties. SF-PCA jointly leverages multiparty homomorphic encryption, interactive protocols, and edge computing to efficiently interleave computations on local cleartext data with operations on collectively encrypted data. SF-PCA obtains results as accurate as non-secure centralized solutions, independently of the data distribution among the parties. It scales linearly or better with the dataset dimensions and with the number of data providers. SF-PCA is more precise than existing approaches that approximate the solution by combining local analysis results, and between 3x and 250x faster than privacy-preserving alternatives based solely on secure multiparty computation or homomorphic encryption. Our work demonstrates the practical applicability of secure and federated PCA on private distributed datasets.

1. Introduction

Principal component analysis (PCA) [1], [2] is an algorithm for analyzing a high-dimensional dataset, represented as a matrix of samples (rows) by features (columns), to uncover a small set of orthogonal directions—principal components (PCs)—that together maximally capture the observed variance among the data samples. Given the ability of PCA to reduce the dimensionality of a dataset while preserving its information content, it is commonly used in many data analysis workflows, including predictive modeling and exploratory data analysis (e.g., clustering and data visualization) [3], [4], [5], [6], [7], [8], [9], [10]. PCA is also a common pre-processing technique in machine learning (ML)

*Equal contribution

pipelines, where the goal is to reduce the number of features to avoid overfitting and improve generalization performance [7], [11], [12], [13], [14]. While more sophisticated non-linear dimension-reduction approaches have been proposed (e.g., based on autoencoders [15], [16]), PCA remains the de-facto standard method for dimension reduction, as it is computationally efficient, theoretically well-understood, and reliably accurate [6], [17].

Many modern applications of PCA involve data from individuals, raising privacy-related challenges that limit the availability of data for such analyses. In the biomedical domain, the high-dimensional nature of biomedical measurements often necessitate the use of PCA to extract key features from personal data, including genetic sequences [9], [18], single-cell transcriptomic data [19], [20], medical images [4], [21] and time-series data [7], [22]. PCA is also commonly used in other domains involving personal data, including quantitative finance [8] and recommender systems [10]. Due to the privacy and security implications, the sharing of personal data in these domains is often prohibited, rendering the data analysis difficult or even impossible. This results in sensitive data remaining siloed in access-controlled repositories and not shared across organizations, which often hinders research, innovation, and routine organizational tasks [23].

Federated privacy-preserving analytics, which aims to facilitate the joint analysis of sensitive data held by multiple parties using privacy-enhancing technologies [24], [25], [26], [27], [28], has emerged as a promising solution to the aforementioned challenges with the potential to overcome regulatory barriers in data sharing [29]. Despite the growing interest, many essential tools for data analysis including the PCA, especially those upstream of widely studied tasks such as model training and inference, have received limited interests and are often omitted from federated workflows. This creates an important gap in secure analytics, potentially undermining their security or utility if one falls back on a non-secure or less-accurate alternative in order to perform the full analysis.

A key challenge in developing a secure federated solution for PCA is that it requires complex and iterative computations (e.g. matrix factorization), which are costly given a large-scale input. These operations are not directly amenable to efficient computation with generic cryptographic techniques [30], [31], [32]. Reflecting this difficulty, many existing federated solutions [33], [34], [35], [36], [37], [38], [39], [40], [41], [42], [43], propose that the data providers (DPs) independently perform an initial dimension reduction on their local data, before they combine their intermediate results and execute the final decomposition on the merged results. This approach, which we refer to as *meta-analysis*, results in a loss of accuracy as it alters the original PCA problem and is prone to overlooking patterns spanning multiple DPs' datasets, especially when the data distributions differ among the DPs. Furthermore, most meta-analysis solutions require the DPs' intermediate results to be revealed to an aggregator server (or to other DPs) hence are not end-to-end secure. Other existing PCA solutions based on secure multiparty computation (SMC) techniques [43], [44], [45], [46] require the entire input data to be securely shared with a few computing servers. With the high communication overhead of SMC, these solutions have difficulty supporting a large number of parties.

In this paper, we propose an efficient and secure system for performing a federated PCA on a distributed dataset, where the data remains protected and locally stored by the respective DPs. Our solution, named SF-PCA (for Secure Federated PCA), executes the randomized PCA (RPCA) algorithm [47], the de facto standard for PCA on large-scale matrices, in a federated manner using a multiparty extension of homomorphic encryption [48]. Contrary to meta-analysis solutions, SF-PCA directly executes a standard PCA algorithm (i.e., RPCA) to achieve state-of-the-art accuracy similar to a centralized analysis, while ensuring end-to-end privacy by protecting even the intermediate results. Unlike SMC solutions, SF-PCA is more communication-efficient and can be used by a large number of DPs. Note that our setting is related to cross-silo federated learning [49], except we do not focus on predictive model training and we use cryptographic techniques to provide end-to-end privacy.

Specifically, SF-PCA is built upon the cryptographic framework of multiparty homomorphic encryption (MHE; see §.3). In MHE, analogous to related works on threshold HE [50], [51], [52], [53], the collective secret (or decryption) key is secret-shared among all the DPs, and the corresponding public key and additional evaluation keys required for homomorphic operations are known by all DPs. This ensures that, while encryption and ciphertext computations can be independently performed by each DP, decrypting ciphertexts requires all DPs to collaborate [48]. MHE's ability to offload certain computations to be locally performed by each party using the cleartext data leads to key performance improvements, as we show in our work. Performing a compute-intensive algorithm like RPCA, which involves sophisticated linear algebra operations (e.g., orthogonalization and eigendecomposition) on input vectors and matrices of a wide range of dimensions, while efficiently working within the constraints of MHE and maximally exploiting its strengths is the key challenge we address in SF-PCA by introducing optimization strategies and efficient MHE linear algebra routines.

Our evaluation demonstrates the practical performance of SF-PCA on six real datasets. For example, SF-PCA securely computes five PCs on the MNIST dataset [54] with 60,000 samples and 760 features, split among six DPs, in 2.22 hours. In the same setting, it obtains the two PCs from a lung cancer dataset [55] with 9,098 patients and 23,724 genomic features in 3.5 hours. SF-PCA scales at most linearly with the input dimensions and with the number of DPs. SF-PCA is one to two orders of magnitude faster than a centralized-HE solution. It is up to ten times faster than existing SMC solutions [45], which scale poorly with the number of DPs. We also show that SF-PCA is highly accurate, resulting in Pearson correlation coefficients of above 0.9 (compared to the ground truth) in all settings, whereas meta-analysis often obtains inaccurate results (e.g., a correlation below 0.75 for both datasets mentioned above). Moreover, SF-PCA executes PCA while ensuring end-to-end data confidentiality as long as one DP is honest, whereas meta-analysis reveals the intermediate results to the aggregator server. Both centralized-HE and the previous SMC solution [45] require an honest third-party to hold the decryption key or to distribute correlated randomness for efficiency, respectively.

In this work, we make the following contributions:

- We propose SF-PCA, a system for an efficient, federated, and end-to-end confidential execution of PCA [47].
- We demonstrate key design strategies underlying the practical performance of SF-PCA, including: (i) maximizing operations on the DPs' cleartext local data by restructuring the computation and (ii) developing efficient linear algebra routines under a consistent vectorized encoding scheme for encrypted matrices to fully utilize the packing and *single-instruction multiple-data* (SIMD) property of MHE without costly encoding conversion.
- We introduce an adaptive approach for choosing both the high-level computational approach for PCA and the low-level MHE routines to maximize efficiency, based on the input dimensions for each computational step.
- We propose efficient MHE-based algorithms for sophisticated linear algebra operations on encrypted matrices, including matrix multiplication, factorization, and orthogonalization, in the federated setting.
- We demonstrate the practical performance of SF-PCA on six real datasets and illustrate its utility for biomedical data analysis. We show that SF-PCA is more scalable than existing solutions for privacy-preserving PCA while producing accurate results comparable to a centralized execution of PCA regardless of the data distribution among the parties.

To the best of our knowledge, SF-PCA is the first system to enable federated PCA in a scalable and end-to-end confidential manner. We note that SF-PCA's optimization strategies and linear algebra building blocks are broadly applicable to the development of secure federated algorithms and thus are of independent interest.

2. Related Work

2.1. Homomorphic Encryption (HE)

We discuss prior works on linear algebra in HE and on distributed HE schemes, two essential components of SF-PCA (§.6).

HE for Linear Algebra.—Multiple works have shown how to optimize matrix-vector multiplications [24], [56] and multiplications between small matrices (i.e., fitting in a single ciphertext) [57], [58], [59], [60]. Multiplication of large encrypted matrices, whose rows do not fit into single ciphertexts, has been less studied. PCA requires multiple types of multiplications involving large matrices of varying dimensions, and efficiently performing these operations under encryption is key to achieving practical performance. SF-PCA jointly leverages a range of matrix multiplication methods whose complexities scale differently with the input dimensions, making an adaptive choice for each computational step in RPCA (§.6.1).

Distributed HE.—When multiple parties use HE to combine their private data, they can either share all of their data encrypted under the same key held by a trusted entity (e.g., in a centralized scheme [61], [62]), or adopt a distributed scheme where no single entity

holds the decryption key. In threshold encryption schemes [63], [64], the encryption key is known to all parties whereas the decryption key is secret-shared among the parties such that a predefined number of them must collaborate to decrypt a ciphertext. In multi-key [65] schemes (including a hybrid with threshold schemes [66]), the parties have their own key pair and can jointly compute on data encrypted under different keys, but the complexity scales with the number of parties. In SF-PCA, we rely on a multiparty HE scheme (MHE) proposed by Mouchet et al. [48], which corresponds to an *s-out-of-s* threshold scheme. This scheme enables local computation with complexity independent of the number of parties and provides a lightweight, interactive protocol to refresh (bootstrap) a ciphertext—a key factor for SF-PCA’s efficiency in contrast to alternative approaches (see §.5).

2.2. Principal Component Analysis (PCA)

Secure Centralized PCA.—Few solutions have been proposed for the secure centralized computation of PCA due to its computational complexity. Pereiral and Aranhall [67] proposed a method for performing PCA on an encrypted dataset using homomorphic encryption (HE). HE-based solutions typically incur a high computational overhead compared to their cleartext counterparts. In addition, they require a costly centralization of the data and have a single point-of-failure, i.e., the holder of the decryption key. In SF-PCA, since the exchanged data are encrypted with a collective key, no single entity can decrypt them, and compute-intensive HE operations (e.g., bootstrapping) are replaced by lightweight interactive protocols. In §.7.6, we compare SF-PCA with an HE-based centralized solution.

Non-Secure Federated PCA.—Solutions that enable PCA on distributed data without privacy protection fall in two main categories: iterative [68], [69], [70], [71] and non-iterative [33], [34], [35], [36], [37], [38], [39], [40], [41], [42], [43]. In the former, the DPs communicate and collaborate in order to perform each step of the algorithm. In the latter, the DPs perform the decomposition locally and then merge their results; we also refer to this approach as *meta-analysis*. Meta-analysis requires less communication but introduces inaccuracies by approximating PCA with two levels of decomposition, i.e., an independent local decomposition by each DP and a global one for the merged results. These solutions typically require that the local data distribution be consistent across DPs to obtain accurate results. In addition, they are not end-to-end secure as they require the DPs’ intermediate results to be revealed to an aggregator server (or to other DPs), representing a single point of failure. Intermediate results have been shown to reveal information about the original data in federated settings, e.g., in PCA [70] and ML [72], [73]. In contrast, SF-PCA implicitly performs RPCA on the joint data without altering the original approach, thus obtaining accurate results independently of the data distribution among the DPs (§.7). It also keeps all the exchanged information secret and does not rely on an aggregator server.

SMC-based PCA.—Several solutions [43], [44], [45], [46] leverage secure multiparty computation (SMC) to perform PCA on data that are secret-shared among a limited number of parties (e.g., three). These solutions require the data to be outsourced to computing parties, incurring a high communication overhead for large datasets. Unlike SMC solutions, SF-PCA can be efficiently used by a large number of parties, and their data are kept locally with a minimal amount of encrypted information exchanged for the PCA computation. In

§.8, we discuss an extension of SF-PCA where SMC techniques are integrated into our system to aid in carrying out non-polynomial function evaluations on small-dimensional inputs.

HE-based PCA.—To our knowledge, Liu et al. [74] proposed the only existing homomorphic encryption (HE)-based solution for federated PCA. However, they rely on an aggregator server that decrypts the aggregated values at each step of the process. Since the intermediate results can reveal information about the parties' local data, these methods are not end-to-end secure. SF-PCA demonstrates that a fully decentralized and end-to-end secure solution for PCA is practically feasible.

Differential Privacy-based PCA.—Solutions based on differential privacy [75], [76], [77] fundamentally differ from SF-PCA in that their goal is to limit the privacy leakage of the intermediate or final results. To achieve this goal, these solutions introduce noise into the computation, making the final results less accurate. Furthermore, analogous to meta-analysis, some of these solutions rely on a local decomposition followed by a global aggregation of results, introducing an approximation error in addition to the noise added for differential privacy. In SF-PCA, no intermediate result is revealed, hence differential privacy is not needed to protect the information exchanged during the algorithm. On the other hand, if the DPs wish to reveal the final PCA result with differential privacy, such guarantee can be added to SF-PCA (§.8).

3. Background

Notation.

Matrices and vectors are denoted by boldface uppercase and lowercase characters, respectively. The i -th row (resp. column) of a matrix $\mathbf{X}^{(a \times b)}$ with a rows and b columns is denoted by $\mathbf{X}[i, :]$ (resp. $\mathbf{X}[:, i]$). The submatrix from (resp. up to but not including) row i and column j is denoted as $\mathbf{X}[i:, j:]$ (resp. $\mathbf{X}[:, :j]$). The i -th element of a vector of b elements $\mathbf{y}^{(b \times 1)}$ is denoted by $\mathbf{y}[i]$. Cleartext data are indicated by a tilde (e.g., $\tilde{\mathbf{X}}$). A matrix multiplication is denoted by \times .

Principal Component Analysis (PCA).

PCA is used to extract the most prominent set of linearly independent directions, i.e., principal components (PCs), that underlie a set of correlated features (columns of a data matrix). The PCs are identified in a descending order of the variance among the data points that each one captures. The PCs can be viewed as the leading eigenvectors of the feature covariance matrix, where the corresponding eigenvalues represent the variance explained. Dimension reduction of the dataset can be achieved by projecting the data points onto the PCs. Formally, PCA takes the matrix $\tilde{\mathbf{A}}^{(n \times m)}$ and outputs the reduced matrix $\tilde{\mathbf{A}}'^{(n \times \psi)}$ obtained from the projection of the input matrix onto its ψ (with $\psi \ll m$) PCs.

Randomized PCA (RPCA)

RPCA [47] is an efficient randomized algorithm for PCA, which lowers the complexity of the matrix decomposition by first reducing the input dimension via random projection [47]. Fig. 1 depicts the workflow of RPCA. It takes as input the matrix $\tilde{\mathbf{A}}$ and a random sketch matrix $\tilde{\mathbf{\Pi}}$ (**Step 1**). We adopt the count-sketch approach [78] for generating the latter, where the elements are drawn from $\{-1,0,1\}$. The columns of the input matrix $\tilde{\mathbf{A}}$ are first mean-centered (**Step 2**); $\tilde{\mathbf{O}}$ denotes the matrix in which each column contains the mean of the corresponding column of $\tilde{\mathbf{A}}$. Next, the input matrix is projected to a lower-dimensional space by multiplying with the sketch matrix (**Step 3**). For improved accuracy [47], the projected matrix $\tilde{\mathbf{P}}$ is recursively multiplied with the covariance matrix $\tilde{\mathbf{A}}^T \tilde{\mathbf{A}}$ for p iterations (**Step 4**). At each iteration, the resulting matrix is orthogonalized using the QR factorization for numerical stability. We denote this step by QR^T , as this algorithm is applied to the rows of the matrix, not columns, in our setting. In **Step 5**, a small symmetric matrix $\tilde{\mathbf{Z}}$ representing the feature covariance in the low-dimensional space is computed by multiplying the result $\tilde{\mathbf{P}}$ of Step 4 on both sides of the covariance matrix. In **Step 6**, the eigenvectors $\tilde{\mathbf{W}}$ of $\tilde{\mathbf{Z}}$ are computed via eigendecomposition (Eigen); we use the QR iteration algorithm with tridiagonalization and implicit shifting of eigenvalues [79] (see §.6.2 for details), which are standard techniques for improving the convergence. RPCA reduces the original problem of factorizing $\tilde{\mathbf{A}} \in \mathbb{R}^{(n \times m)}$ to decomposing the tiny, constant-size matrix $\tilde{\mathbf{Z}} \in \mathbb{R}^{(\rho \times \rho)}$, where $\rho = \psi + \alpha$ with ψ the desired number of principal components and α an oversampling parameter. The latter is used to increase the accuracy of the algorithm [47]. In **Step 7**, it reconstructs the eigenvectors in the original space (i.e., the PCs $\tilde{\mathbf{W}}$) and finally projects the data points of $\tilde{\mathbf{A}}$ onto the PCs in **Step 8** to construct the output.

Multiparty Homomorphic Encryption (MHE).

To securely perform PCA across distributed datasets, we rely on a multiparty (or distributed) fully-homomorphic encryption scheme [48] in which the secret key sk is shared among the parties via a secret-sharing scheme, whereas the corresponding collective public key pk is known to all of them. As a result, each party can independently compute on ciphertexts encrypted under pk , but all parties have to collaborate to decrypt a ciphertext.

Mouchet et al. [48] showed how to adapt ring-learning-with-errors-based homomorphic encryption schemes [61], [62], [80] to the multiparty setting. In SF-PCA, we instantiate the multiparty scheme with the Cheon-Kim-Kim-Song (CKKS) cryptosystem [61]. CKKS is a homomorphic encryption scheme that enables approximate arithmetic over $\mathbb{C}^{\mathcal{N}/2}$; the plaintext and ciphertext spaces share the same domain $R_{Q_L} = \mathbb{Z}_{Q_L}[X]/(X^{\mathcal{N}} + 1)$, with $Q_L = \prod_0^L q_i$ in our case and \mathcal{N} a power of 2. Both plaintexts and ciphertexts are represented by polynomials of degree up to $\mathcal{N} - 1$ (with \mathcal{N} coefficients) in this domain, each encoding a vector of up to $t = \mathcal{N}/2$ floating-point values. Any operation is SIMD, i.e., simultaneously performed on all encoded values. CKKS's security is based on the ring learning with errors (RLWE) problem [80] and some noise is added directly in the least significant bits of the encrypted values. Mouchet et al. [48] have shown that the distributed protocols (described below) introduce only additive noise, linear in the number of DPs. To limit the noise

growth during homomorphic operations in SF-PCA, we leverage general scale-management techniques for CKKS [81], [82], [83]. We refer to Appendix A for cryptoscheme details.

Main MHE Operations.—The DPs each have a public key pk_i and the corresponding secret key sk_i (with $\{sk_i\}$ the set of all DPs' secret keys) and can collectively execute the following operations. We denote a collectively encrypted vector by c and a plaintext vector by \tilde{p} . Symbols are summarized in Tab. 3 (Appendix B).

- $pk, evks \leftarrow \text{DKeyGen}(\{sk_i\})$ generates the collective public key pk and evaluation keys $evks$, which are required for ciphertext transformations such as rotations. The DPs aggregate the local shares of keys (randomly generated based on a public source of randomness) to obtain public collective keys [48].
- $c \leftarrow \text{DBootstrap}(c', \{sk_i\})$ collectively refreshes a ciphertext to obtain a fresh encryption. This operation is required after every λ multiplications to ensure a correct decryption.
- $c_{pk'} \leftarrow \text{DKeySwitch}(c, pk', \{sk_i\})$ changes the encryption of a ciphertext c from the public key pk to another public key pk' , without decrypting the ciphertext. The collective decryption is a special case of this operation (i.e., $\text{DKeySwitch}(c, \emptyset, \{sk_i\})$). To prevent information leakage upon decryption [84], a fresh noise with a variance larger than that of the ciphertext is added before decryption [48], [84], [85], [86].

Each DP can independently encrypt, and perform the following operations listed in order of increasing computational complexity (Tab. 2):

- $c_{pk} \in R_{QL}^2 \leftarrow \text{Enc}(pk, \tilde{p})$ with a plaintext vector \tilde{p} , such that $\text{DKeySwitch}(c_{pk}, \emptyset, \{sk_i\}) \approx \tilde{p}$.
- $c'' = c + c'$, addition of encrypted vectors.
- $c' = c \cdot \tilde{p}$, element-wise multiplication of an encrypted vector and a cleartext vector. The result needs to be rescaled to maintain ciphertext scale.
- $c'' = c \cdot c'$, element-wise multiplication of two encrypted vectors. The result needs to be relinearized and rescaled to maintain ciphertext size and scale.
- $c' = \text{Rot}_y(c, evks)$, cyclic rotation of length y to the left (to the right if y is negative) on the encrypted vector c .
- $c'' = c \cdot c'$, dot product of two encrypted vectors. The result is encoded in the first position of a one-hot encoded vector c'' .
- $c' = \text{Dup}_y(c)$, duplication of the first element of c to the first y positions of c' with $\lceil \log_2(y) \rceil$ rotations and additions.

4. SF-PCA System and Security Models

Our system model is illustrated in Fig. 2. We consider a cleartext dataset, represented as a matrix $\tilde{A} \in \mathbb{R}^{(n \times m)}$, that is horizontally split among a set of interconnected data providers DP_1, \dots, DP_s such that each DP_i has $\tilde{A}_i \in \mathbb{R}^{(n_i \times m)}$ with $\sum_i n_i = n$. The number of data samples held by each DP (i.e., n_i) is considered public. We discuss the vertically partitioned case in §.8.

SF-PCA enables the DPs to collaboratively execute a randomized PCA on their joint data. In the end, each DP obtains ψ collectively encrypted PCs, on which each DP can locally project its data. If required by the application, each DP's projected data (encrypted under the collective key) can be collectively switched (DKeySwitch, §.3) to each DP's public key pk_i to be locally decrypted. Similarly, the PCs can be collectively decrypted and shared among the DPs.

We adopt the semi-honest model, where the DPs follow the protocol as specified, but might try to infer information about another DP's data, potentially colluding with other DPs. We require that the DPs' data and all intermediate results remain confidential. In other words, SF-PCA provides *input confidentiality*, i.e., no DP is able to learn any information about any other DP's local data other than what it can infer from the final output of PCA (e.g., its projected local data). We require that this property holds as long as one DP remains honest and does not collude with others.

5. SF-PCA Protocol Design

We introduce an end-to-end confidential and federated approach to execute a RPCA (§.3) jointly over s DPs holding their local data. At each step of the PCA execution, the DPs collectively compute encrypted global intermediate results through interactive protocols that combine the results of local computation on each DP's cleartext data. The intermediate results remain encrypted under the DPs' collective key and are never revealed. While our system's ability to leverage local cleartext computation opens the door to efficient multiparty algorithms, a careful algorithmic design is still necessary for developing a practical PCA protocol.

Leveraging existing approaches for secure computation (e.g., HE or SMC), the DPs could outsource their encrypted (or secret-shared) data to one or multiple computing parties to jointly perform the PCA. However, the communication overhead of sharing the entire dataset as well as the computational burden of performing complex computations (e.g., multiplication and factorization of matrices) on the pooled dataset render these solutions impractical for large-scale datasets. Note that the repeated matrix multiplications are challenging to perform efficiently under HE due to the costly bootstrapping procedure. SF-PCA addresses these challenges by introducing efficient MHE-based protocols based on a federated approach to joint computation. We compare SF-PCA's performance with existing approaches in §.7.

5.1. Key Strategies for Accuracy and Efficiency

In RPCA (Fig. 1), many matrix multiplications involving the input matrix $\widetilde{\mathbf{A}}^{(n \times m)}$ (or its covariance matrix) largely determine the protocol's complexity. These multiplications are interspersed with sophisticated linear algebra transformations, such as the QR factorization invoked at the end of each power iteration (Step 4 in Fig. 1) and eigendecomposition (Eigen in Step 6), which view the matrix as a set of row (or column) vectors and apply vector-level operations. Below, we explain our strategies to carry out these computations efficiently while maintaining the accuracy of results.

Obtaining Accurate Results by Emulating Centralized PCA.—Existing federated approaches to PCA that combine the results independently obtained by the DPs (e.g., meta-analysis), are prone to errors introduced by differences in data distribution among the DPs. In SF-PCA, we avoid this pitfall by securely combining the intermediate results at each step of the protocol (via collective aggregation Ξ in Alg.1) to emulate a centralized analysis, thus obtaining the same PCs regardless of how the data are split (§.7.7).

Efficient Edge-Computing on Local Cleartext Data.—Working with an encrypted form of the entire input matrix ($\widetilde{\mathbf{A}}$ in Fig. 1) would require the DPs to transfer large amounts of data (e.g., for centralized HE or secret sharing) or to perform costly ciphertext operations on large matrices, both of which become impractical for large-scale datasets. In SF-PCA (Alg.1), the DPs jointly perform the PCA without encrypting or exchanging the input data. They collaborate instead by computing on their local cleartext data (i.e., the sub-matrix $\widetilde{\mathbf{A}}_i$) and exchanging only low-dimensional and aggregate-level encrypted information. This enables the DPs to minimize communication and maximize the use of low-cost MHE operations involving the cleartext data (e.g., with our default parameters, cleartext-ciphertext multiplication is eight times faster than a ciphertext-ciphertext multiplication; Tab. 2). We also modify the RPCA computation to use only the cleartext input throughout the workflow. For example, instead of directly constructing a mean-centered input matrix (Step 2 in Fig. 1), which needs to be encrypted due to the means being private, SF-PCA keeps each local matrix $\widetilde{\mathbf{A}}_i$ in cleartext and associates with it an encrypted mean vector σ to correct for mean shifts in subsequent steps (see Step 2 in Alg.1). This enables a key optimization for the matrix multiplications in Steps 3–5, 7 and 8 (Alg.1), where the cleartext matrix $\widetilde{\mathbf{A}}_i$ is pre-transformed to minimize costly ciphertext operations such as rotations in later steps. In §.6.1.3, we show how to efficiently multiply an encrypted matrix with another containing only duplicated rows (or columns), which is used for lazy mean correction in Steps 4, 5, 7 and 8.

Adaptive Selection of Computational Routines based on Data Dimensions.—In practice, PCA is applied to datasets whose dimensions vary greatly depending on the application, e.g., from tens of features in small predictive modeling tasks to tens of thousands of features in genomic studies (§.7). To achieve practical performance in a wide range of settings, we propose an adaptive approach for optimizing the computational routines based on the input dimensions. In Alg.1, we introduce two different workflows for performing RPCA: *Precomp* and *Seq*. In *Precomp*, the encrypted covariance matrix \mathbf{G}

is precomputed in the beginning of Step 4 and reused, such that most of the following operations scale primarily with the number of features m . In *Seq*, $\tilde{\mathbf{A}}$ is kept in cleartext and used for matrix multiplications, which is more efficient than using \mathbf{G} , but now the computation scales with both m and the number of samples n . In addition, in §.6.1, we describe several matrix multiplication methods, each of which scales differently with the input dimensions; SF-PCA selects the best approach for each step in its workflow. Similarly, in §.6.2, we introduce two approaches for performing the QR factorization on an encrypted matrix (QR in Step 4), with different complexities depending on the input dimensions.

Optimized Data Encoding for Linear Algebra on Encrypted Matrices.—The secure execution of RPCA requires that the DPs iteratively perform various matrix operations on encrypted data, including multiplication and factorization. For example, the QR factorization, which is repeatedly executed in-between matrix multiplications (in Steps 4, 6, and 7 of RPCA; Fig. 1), is performed over the rows of an encrypted matrix in SF-PCA. Selecting a row in a matrix of m columns, where the columns are individually packed in ciphertexts, would require m homomorphic multiplications, additions, and rotations; in contrast, row selection incurs no cost when the matrix is row-wise encoded. In fact, the overwhelming cost of transforming encrypted matrices from one encoding to another would make our system impractical. We therefore adopt a consistent vectorized encoding scheme throughout the algorithm to represent encrypted matrices and tailor the operations to efficiently work with this format without costly conversions. This also allows SF-PCA to fully utilize the packing and SIMD properties of MHE thus minimizing its overall computation and communication costs.

Selective Bootstrapping to Minimize Communication.—After a certain number of multiplications, a ciphertext needs to be bootstrapped (DBootstrap, §.3) to restore its capacity for computation. In SF-PCA, this is a collective operation, which is computationally lightweight in contrast to its centralized equivalent, but requires the ciphertext to be exchanged among all DPs. To further minimize this communication overhead, we restrict the invocation of DBootstrap to places where an intermediate result is already globally synced and of a small dimension (e.g., during QR factorization in Steps 4 and 7 in Alg. 1; see §.6.2), while flexibly allowing a ciphertext to be bootstrapped even if some multiplication capacity remains.

5.2. Workflow Details

We describe the workflow of SF-PCA from the point of view of DP, in Alg. 1. Recall that the DPs aim to compute ψ encrypted PCs (rows of matrix \mathbf{W}) on their joint data. RPCA identifies $\rho = \psi + \alpha$ components with a small oversampling parameter α for improved accuracy. In addition to Alg. 1, we show in Fig. 8 in the Appendix how the matrix dimensions evolve in SF-PCA's workflow. The DPs interact by aggregating (represented by Ξ) encrypted matrices and broadcasting the encrypted result to all DPs.

Step 1: Setup.—Each DP _{i} holds $\tilde{\mathbf{A}}^{(n_i \times m)}$, a submatrix of the global input matrix $\tilde{\mathbf{A}}^{(n \times m)}$. The DPs generate the required public keys (DKeyGen, §.3) and agree on the PCA

parameters, including: the number of power iterations p , the number of QR iterations w for eigendecomposition, the desired number of PCs ψ , the oversampling parameter α (resulting in the number of components $\rho = \psi + \alpha$ for RPCA), and a public random sketch matrix $\tilde{\mathbf{I}}^{(\rho \times n)}$ (e.g., generated from a shared seed). In addition, the DPs together decide the specifics of certain computational steps in SF-PCA, such as the approximation intervals for non-linear operations (§.6.3) and the method of choice for costly linear algebraic operations (matrix multiplication and transformations; see §.6), taking the input dimensions into account to maximize performance. All the parameters introduced in this step are considered public. Note that the procedure to agree upon the parameters is orthogonal to SF-PCA; e.g., the DP initiating the collaboration could propose the parameters.

Step 2: Mean Calculation.—The DPs compute the encrypted vector $\mathbf{o}^{(1 \times m)}$ of column averages of the input matrix $\tilde{\mathbf{A}}^{(n \times m)}$ by securely aggregating their local column sums divided by n , encrypted under the collective public key.

Algorithm 1 - SF-PCA (from DP_i point of view)

Private Input: $\tilde{A}_i^{(n_i \times m)}$, **Public Input:** $\rho, \psi, \alpha, \tilde{\Pi}^{(\rho \times n)}$, $m, n, \{n_1, \dots, n_s\}, p, w$, approx. params., $\rho = \psi + \alpha \ll m, n$

Private Output: $((A'_i)^T)^{(\psi \times n_i)}$ & $W^{(\psi \times m)}$

Notation: Ξ is collective aggr., vectors are replicated to match matrix dims. A tilde indicates cleartext.

Step 1: $pk, evks \leftarrow \text{DKeyGen}(sk_i)$

Step 2: $\tilde{o}_i^{(1 \times m)} = [\frac{1}{n}(\sum \tilde{A}_i[:,0]), \dots, \frac{1}{n}(\sum \tilde{A}_i[:,m-1])]$
 $o^{(1 \times m)} = \Xi(\text{Enc}(pk, \tilde{o}_i^{(1 \times m)}))$

Step 3: $P^{(\rho \times m)} = \Xi(\text{Enc}(pk, \tilde{\Pi}_i^{(\rho \times n_i)} \times \tilde{A}_i - \tilde{\Pi}^{(\rho \times n)} \times \tilde{o}))$

Step 4:

<p>Precomp</p> $G_i^{(m \times m)} = \tilde{A}_i^T \times \tilde{A}_i - \tilde{A}_i \times o - \tilde{A}_i^T \times o^T + o^T \times o$ <p>for $j=0, \dots, p-1$ do $P^{(\rho \times m)} = \text{QR}^T(\Xi(P \times G_i))$ end for</p>	or	<p>Seq</p> <p>for $j=0, \dots, p-1$ do $R_i^{(\rho \times n_i)} = P \times \tilde{A}_i^T - P \times o^T$</p> <p>DQR</p> $R_i^{(\rho \times n_i)} = \text{DQR}^T(R_i)$ $P^{(\rho \times m)} = \Xi(R_i \times \tilde{A}_i - R_i \times o)$ <p style="text-align: center;">or</p> <p>QR</p> $P^{(\rho \times m)} = \text{QR}^T(\Xi(R_i \times \tilde{A}_i - R_i \times o))$ <p>end for</p>
---	----	---

Step 5:

<p>Precomp</p> $P_i^{(\rho \times m)} = P \times G_i$ $Z_i^{(\rho \times \rho)} = \Xi(P_i \times P_i^T)$	or	<p>Seq</p> $Z_i^{(\rho \times n_i)} = P \times \tilde{A}_i^T - P \times o^T$ $Z_i^{(\rho \times \rho)} = \Xi(Z_i \times Z_i^T)$
---	----	--

Step 6: $U^{(\psi \times \rho)} = \text{Eigen}(Z^{(\rho \times \rho)}, w)$

Step 7:

<p>Precomp</p> $W^{(\psi \times m)} = \text{QR}^T(\Xi(U \times P_i))$	or	<p>Seq</p> $Y_i^{(\psi \times n_i)} = U \times Z_i^T$ $W^{(\psi \times m)} = \text{QR}^T(\Xi(Y_i \times \tilde{A}_i - Y_i \times o))$
--	----	--

Step 8: $((A'_i)^T)^{(\psi \times n_i)} = W \times \tilde{A}_i^T - W \times o^T$

Step 3: Random Projection.— DP_i projects \tilde{A}_i to a subspace of ρ dimensions using the public sketch matrix $\tilde{\Pi}^{(\rho \times n)}$. DP_i locally computes the product of \tilde{A}_i and the corresponding submatrix $\tilde{\Pi}_i^{(\rho \times n_i)}$ of $\tilde{\Pi}^{(\rho \times n)}$ to obtain its local sketch in cleartext. The result is then encrypted and aggregated among all DPs to obtain the encrypted sketch P of the global matrix \tilde{A} .

Step 4: Power Iterations.—The sketch of the input matrix obtained in the previous step is repeatedly multiplied with the input matrix to increase the spectral gap between the top eigenvectors of interest and the rest [47]. We execute this step differently depending on the input dimensions for optimized performance; the two approaches considered by SF-PCA are described below. Notably, in both approaches, we leverage the fact that the cost of cleartext operations is almost negligible compared to that of HE to optimize the computation.

Approach 1: Precompute & Reuse (Precomp): Each DP_i precomputes the covariance matrix $G_i^{(m \times m)}$ once and reuses it in every iteration for multiplying with P . Note that G_i needs to be encrypted due to the mean-centering operation using the encrypted global column means $o^{(1 \times m)}$. SF-PCA's optimized matrix multiplication routines between an encrypted and a cleartext matrix (§.6.1) minimize the computation involving the encrypted matrix by precomputing certain transformations of the cleartext matrix at a negligible cost. To efficiently apply these methods to the encrypted G_i , we obtain the transformations of G_i by transforming the cleartext \tilde{A}_i and \tilde{A}_i^T before multiplying them.

Approach 2: Sequentially Multiply (Seq): The DPs sequentially multiply P by the cleartext matrix \tilde{A}_i (and its transpose) on the fly. The covariance matrix is never explicitly constructed. To keep the input matrix \tilde{A}_i in cleartext, SF-PCA performs the mean-centering of \tilde{A}_i in a *lazy* manner (*lazy mean-centering*): instead of subtracting the encrypted vector $o^{(1 \times m)}$ from each row of \tilde{A}_i , which would transform the whole input matrix into an encrypted matrix, multiplication is performed using the original cleartext \tilde{A}_i and the resulting matrix is corrected to account for the mean shift. More precisely, we multiply encrypted P with mean-centered \tilde{A}_i in three efficient steps: (1) multiply P with the cleartext matrix \tilde{A}_i , (2) compute the inner product between each row of P and o (see §.6.1.3), and (3) subtract each inner product value from all elements in the corresponding row of the matrix from (1). We observe that *Precomp* requires fewer multiplications per power iteration and its computation cost is mostly independent of n . *Seq* requires more operations but maximizes cleartext operations by reusing the cleartext matrix \tilde{A} . We compare the performance of both approaches in §.7.

In each iteration, a QR factorization (QR^T ; Alg. 3) is applied to either the aggregated matrix $P^{(\rho \times m)}$, in both approaches, or the intermediate $(P \times \tilde{A}_i)^{(\rho \times n_i)}$ in *Seq*. In the latter, the factorization is optionally performed using a new interactive protocol DQR^T , when the computational speedup of each DP computing on a matrix with n_i columns vs. one DP computing on an aggregated matrix with m columns exceeds the additional communication cost, i.e., when $n_i \ll m$ (see §.6.2).

Step 5: Reduction.—In the *Precomp* approach, the matrix P resulting from Step 4 is transformed to a small symmetric matrix Z by multiplying the covariance matrix G_i from both sides. In the *Seq* approach, this is performed by using the cleartext matrix \tilde{A}_i^T and then by multiplying the result by its transpose. As in Step 4, SF-PCA employs *lazy mean-centering* for this step.

Step 6: Eigendecomposition.—The eigendecomposition (introduced in §.3) is executed on the encrypted matrix Z . We detail our MHE-based algorithm for this step in Alg. 4. It requires the iterative execution of QR^T and matrix multiplications.

Step 7: Reconstruction.—The PCs (rows of \mathbf{W}) are computed by multiplying the eigenvectors from Step 6 with the approximated subspace from Step 3, followed by a final round of power iteration and orthogonalization (\mathbf{QR}^T) for numerical stability.

Step 8: Projection.—Each DP_{*i*} projects its local cleartext data $\tilde{\mathbf{A}}_i^T$ onto the collectively encrypted PCs in \mathbf{W} to obtain their projected data \mathbf{A}'_i , which is also encrypted under the collective public key. If required by the application, by using DKeySwitch, the PCs and/or the DP_{*i*}'s projected data can be collectively decrypted or re-encrypted under the public keys of specific entities to grant controlled access to the decrypted results.

6. Optimized Routines for Linear Algebra and Non-Polynomial Functions on Encrypted Data

We describe how SF-PCA efficiently executes matrix multiplications, sophisticated linear-algebra transformations and non-polynomial function evaluations on encrypted data. Although the methods in this section can also be employed in the centralized setting, we note that the adaptive use of matrix multiplication routines and the higher-level protocols for matrix transformations (e.g., QR factorization) are optimized while accounting for the unique properties of MHE, e.g., the availability of local cleartext data and a lightweight interactive bootstrapping routine, both of which alter the tradeoff between different computational strategies and present new ways to optimize the algorithm. Our secure federated routines may be of independent interest for other applications.

6.1. Matrix Multiplications

Encrypted matrix multiplications are frequently invoked in SF-PCA's workflow and hence are a key determinant of its performance. As outlined in Alg. 1, we introduce two high-level algorithmic workflows—*Precomp* and *Seq*—for executing RPCA. Both approaches involve different types of multiplications over matrices of varying dimensions, motivating our adaptive strategy for choosing the most efficient routine for each computational step in SF-PCA among a range of multiplication methods.

6.1.1. Adaptive Strategy.—We identify two main types of matrix multiplications in Alg. 1: (i) *unbalanced multiplications* between a large encrypted matrix and a large cleartext (or pre-transformed encrypted) matrix in Steps 4, 7 and 8, with the key property that operations are cheap on one matrix (cleartext) and expensive on the other (ciphertext); and (ii) *duplicated-vector multiplications*, referring to multiplications between a large encrypted matrix and another encrypted matrix whose rows (or columns) are identical (e.g., corresponding to the encrypted mean vector \mathbf{o}).

In §.6.1.2, we detail three different approaches (M1, M2, and M3) for unbalanced multiplications, each with a complexity that scales differently with the input dimensions. We denote by $\zeta_M^*(a, b, c)$, the function that takes the three input dimensions for multiplying $\mathbf{M}^{(a \times b)}$ and $\tilde{\mathbf{N}}^{(b \times c)}$ matrices and outputs the cost associated with the most efficient multiplication routine. The cost we compare is a weighted sum of the multiplication and

rotation invocation counts, where the weights are determined by the estimated runtime per operation in the given computational environment. The cost of a cleartext-ciphertext multiplication is set to be around 8 times lower than that of a ciphertext-ciphertext multiplication according to our estimates (Tab. 2).

We identify the matrix multiplication costs of *Precomp* and *Seq* for a single iteration as $\zeta_M^*(\rho, m, m)$ and $\zeta_M^*(\rho, m, n_{\max}) + \zeta_M^*(\rho, n_{\max}, m)$, respectively, where ρ is the number of reduced dimensions in RPCA, m is the number of input features, and $n_{\max} = \max_i(n_i)$ represents the largest number of samples locally held by the DPs. We consider the worst-case complexity as the overall runtime being as fast as the slowest DP. In addition, for both approaches, we incorporate the cost of lazy mean-centering when comparing the overall cost; *Precomp* requires $3b$ Mults and b Rots, whereas *Seq* requires two *duplicated-vector multiplications*, which we detail in §.6.1.3. We further compare these approaches in §.7.

6.1.2. Unbalanced Multiplications.—We describe the HE implementations of three matrix multiplication strategies: Dot-Product Method (M1), Element-Duplication Method (M2), and Diagonal Method (M3; adapted from Jiang et al. [59]). We jointly consider these three methods because their costs scale differently with the input dimensions, enabling SF-PCA to optimize its performance in a wide range of scenarios. For each method, we show its cost in terms of the invocations of ciphertext rotations (Rots) and multiplications (Mults) for multiplying a pair of $a \times b$ and $b \times c$ matrices. The cost of cleartext operations is negligible. To simplify the computational complexity analysis, we assume that b and c are powers of two without loss of generality. We denote by t the ciphertext capacity, i.e., the number of values that can be packed in a ciphertext. Due to SF-PCA’s vectorized encoding, the inner dimension b reduces to a small constant $\lceil \frac{b}{t} \rceil$ in terms of the number of ciphertext operations.

Dot-Product Method (M1): Each element of \mathbf{R} is obtained from the dot product (\bullet) between a row of \mathbf{M} and a column of $\widetilde{\mathbf{N}}$ (Line 4 in M1). The result of the dot-product is moved to position j (0-based) by masking and rotating the vector by j positions to the right (i.e., $\text{Rot}_{(-j)}$; see §.3). In SF-PCA, this method is used (in Step 5 in Alg. 1) to multiply an encrypted matrix by its transpose without any additional transformation, since the encrypted rows of \mathbf{M} can be directly used as the columns of \mathbf{N} . The multiplication and rotation costs mainly depend on the outer dimensions.

M1: Dot-Product Method

Input: Encrypted $\mathbf{M}^{(a \times b)}$ and cleartext (indicated by a tilde) $\widetilde{\mathbf{N}}^{(b \times c)}$.

Output: Encrypted $\mathbf{R}^{(a \times c)} = \mathbf{M} \times \widetilde{\mathbf{N}}$

Cost: $(\lceil \frac{b}{t} \rceil + 1) \cdot ac$ Mults and $ac \cdot \lceil \frac{b}{t} \rceil \cdot \log_2(t)$ Rots

$$1: \quad \mathbf{R} \leftarrow \mathbf{0}^{(a \times c)}$$

```

2:   for  $i = 0, \dots, a - 1$  do
3:       for  $j = 0, \dots, c - 1$  do
4:            $\mathbf{R}[i, :] \leftarrow \mathbf{R}[i, :] + \text{Rot}_{(-j)}(\mathbf{M}[i, :] \bullet \widetilde{\mathbf{N}}[:, j])$ 
5:       end for
6:   end for

```

Element-Duplication Method (M2).: This method avoids the computation of pairwise dot products (used in M1) by duplicating each element of \mathbf{M} to construct a vector of length c and by multiplying this vector (element-wise) with each row of $\widetilde{\mathbf{N}}$ (Line 4 in M2). The results are aggregated to obtain \mathbf{R} . This method's cost depends mostly on the left matrix dimensions.

M2: Element-Duplication Method

Input: Encrypted $\mathbf{M}^{(a \times b)}$ and cleartext (indicated by a tilde) $\widetilde{\mathbf{N}}^{(b \times c)}$

Output: Encrypted $\mathbf{R}^{(a \times c)} = \mathbf{M} \times \widetilde{\mathbf{N}}$

Cost: $\lceil \frac{b}{t} \rceil \cdot ab$ Mults and $ab \cdot \log_2(\min\{c, t\})$ Rots

```

1:    $\mathbf{R} \leftarrow \mathbf{0}^{(a \times c)}$ 
2:   for  $i = 0, \dots, a - 1$  do
3:       for  $j = 0, \dots, b - 1$  do
4:            $\mathbf{R}[i, :] \leftarrow \mathbf{R}[i, :] + (\text{Dup}_c(\mathbf{M}[i, j]) \cdot \widetilde{\mathbf{N}}[j, :])$ 
5:       end for
6:   end for

```

Diagonal Method (M3).: This approach is based on the technique of Jiang et al. [59], which we adapt to large-scale matrices that cannot be packed in a single ciphertext. This method transforms the cleartext matrix (by rotating its columns) such that one of its rows corresponds to the diagonal of the original matrix (Line 2 in M3). The rows of the encrypted \mathbf{M} are then rotated (Line 8) before being multiplied with the transformed rows of $\widetilde{\mathbf{N}}$ at each iteration along the common dimension b (Line 9). We use the baby-step giant-step approach [87] to reduce the number of rotations on the rows of \mathbf{M} from b to $2\sqrt{b}$ by storing the intermediate results in three-dimensional tensors (i.e., \mathbf{M}' and \mathbf{R}') (Lines 8 and 9), introducing a tradeoff between computation and memory usage (see §.7.4). The intermediate results are then aligned and aggregated in the final matrix \mathbf{R} (Line 14). The rows of \mathbf{M} are duplicated or truncated to have c elements ($\text{Len}_c(\cdot)$) before the multiplication. This method's cost also depends mostly on the dimension of the left matrix but, contrary to M2, its number of rotations scales with the square root of the inner dimension times the number of packed ciphertexts along the same dimension.

M3: Diagonal Method

Input: Encrypted $M^{(a \times b)}$ and cleartext (indicated by a tilde) $\widetilde{N}^{(b \times c)}$

Output: Encrypted $R^{(a \times c)} = M \times \widetilde{N}$

Cost: $\lceil \frac{b}{t} \rceil \cdot ab$ Mults and $\lceil \frac{b}{t} \rceil \cdot (\lceil \frac{c}{b} \rceil + 2a \cdot \lceil \sqrt{b} \rceil)$ Rots

- 1: **for** $i = 0, \dots, c - 1$ **do**
- 2: $\widetilde{N}[:, i] \leftarrow \text{Rot}_t(\widetilde{N}[:, i])$
- 3: **end for**
- 4: $M' \leftarrow \mathbf{0}^{(a \times \lceil \sqrt{b} \rceil \times c)}, R' \leftarrow \mathbf{0}^{(a \times \lceil \sqrt{b} \rceil \times c)}$
- 5: **for** $i = 0, \dots, b - 1$ **do**
- 6: $y \leftarrow i \bmod \lceil \sqrt{b} \rceil, g \leftarrow \lfloor i / \lceil \sqrt{b} \rceil \rfloor$
- 7: **for** $j = 0, \dots, a - 1$ **do**
- 8: **if** $M'[j, y, :] = \emptyset$ **then** $M'[j, y, :] \leftarrow \text{Len}_t(\text{Rot}_t(M[j, :]))$
- 9: $R'[j, g, :] \leftarrow R'[j, g, :] + M'[j, y, :] \cdot \widetilde{N}[i \bmod m, :]$
- 10: **end for**
- 11: **end for**
- 12: **for** $i = 0, \dots, a - 1$ **do**
- 13: **for** $l = 0, \dots, \lceil \sqrt{b} \rceil - 1$ **do**
- 14: $R[i, :] \leftarrow R[i, :] + \text{Rot}_{t \cdot \lceil \sqrt{b} \rceil}(R'[i, l, :])$
- 15: **end for**
- 16: **end for**

6.1.3. Duplicated-Vector Multiplications.—This method addresses a special setting where we multiply an encrypted matrix M with another encrypted matrix Γ whose rows (*Case 1*) or columns (*Case 2*) are identical vectors μ . This setting frequently arises in SF-PCA for the *lazy mean-centering* operations (i.e., all operations involving \mathbf{o} in §.5.2). Our method accounts for this redundancy in the matrix to minimize the number of rotations on both encrypted matrices.

M4: Vector-Duplication Method

Input: Encrypted $M^{(a \times b)}$ and encrypted $\Gamma^{(b \times b)}$

where $\Gamma = 1 \times \mu^T$ (*Case 1*) or $\Gamma = \mu \times 1^T$ (*Case 2*)

Output: Encrypted $R^{(a \times b)} = M \times \Gamma$

Cost: $\lceil \frac{b}{t} \rceil \cdot a$ Mults and $\lceil \frac{b}{t} \rceil \cdot 2a \cdot \log_2(\min\{b, t\})$ Rots

- 1: $R \leftarrow \mathbf{0}^{(a \times b)}$

```

2:  for  $i = 0, \dots, \text{ado}$ 
3:    if Case 1 then  $\mathbf{R}[i, :] \leftarrow \text{Dup}_b(\mathbf{M}[i, :] \bullet 1) \cdot \boldsymbol{\mu}$ 
4:    if Case 2 then  $\mathbf{R}[i, :] \leftarrow \text{Dup}_b(\mathbf{M}[i, :] \bullet \boldsymbol{\mu})$ 
5:  end for

```

6.1.4. Further Optimizations.—We note that all multiplication methods are parallelizable at the row level. Each multiplication of a ciphertext is followed by a rescale and a relinearization operation (the result of a multiplication with a plaintext only needs to be rescaled, see §.3). When the results of several multiplications need to be aggregated, we defer the rescale and relinearization operations until *after* the aggregation step so they can be executed once overall, rather than for every multiplication. Because these operations account for between 52% and 75% of the multiplication time (Tab. 2) and SF-PCA heavily relies on matrix multiplications, this optimization considerably improves SF-PCA’s overall performance. For example, it reduces the runtime of multiplying an encrypted $\mathbf{M}^{(8 \times 2^8)}$ with a cleartext $\tilde{\mathbf{N}}^{(2^8 \times 2^8)}$ with M3 from 24.6 to 3.8 seconds; the improvement is expected to be greater for larger matrices.

6.2. Matrix Transformations and Factorizations

We introduce new routines for executing sophisticated linear algebra operations required by the PCA on encrypted matrices and vectors. We begin with the Householder transformation [88], a key building block in other matrix transformations such as QR^T and Eigen, which we subsequently describe. We also present a new algorithm, DQR^T , for executing a QR factorization on a matrix that is distributed among multiple parties. Note that all methods except DQR^T require communication only for bootstrapping (DBootstrap; §.3), which has a negligible computation cost. The reported communication costs are thus measured by our optimized number of invocations of bootstrapping on a single ciphertext.

Algorithm 2 - Encrypted Householder Vector (**HH**)

Input: Encrypted $\mathbf{v}^{(h \times 1)}$

Output: Encrypted $\mathbf{v}'^{(h \times 1)}$, such that $\mathbf{H} = \tilde{\mathbf{I}}^{(h \times h)} - 2\mathbf{v}' \times \mathbf{v}'^T$ ensures $\mathbf{H} \times \mathbf{v}$ all zeros except the first coordinate

Comp. Cost: $3 \cdot l(d) + 6$ Mults and $2 \cdot \log_2(h)$ Rots; $l(d)$ defined in text

Comm. Cost: $(\frac{5 + 3(\lceil 1 + \log_2(d) \rceil)}{\lambda}) \cdot \lceil \frac{h}{t} \rceil$ Ciphertexts

1: $\mathbf{v}^2 \leftarrow \mathbf{v} \cdot \mathbf{v}$	6: $\mathbf{u} \leftarrow \mathbf{v}$
2: $\ \mathbf{v}\ ^2 \leftarrow \mathbf{v}^2 \bullet 1$	7: $\mathbf{u}[0] \leftarrow \delta + \mathbf{v}[0]$
3: $\ \mathbf{v}\ \leftarrow \sqrt{\ \mathbf{v}\ ^2}$	8: $\mathbf{u}^2 \leftarrow \mathbf{u} \cdot \mathbf{u}$
4: $\delta \leftarrow \mathbf{v}[0] / \sqrt{\mathbf{v}[0]^2}$	9: $\mathbf{k} \leftarrow \mathbf{u}^2[0] + (\ \mathbf{v}\ ^2 - \mathbf{v}^2[0])$

$$5: \delta \leftarrow \delta \cdot \| \mathbf{v} \|$$

$$10: \mathbf{k}' \leftarrow \text{Dup}_n(1/\sqrt{\mathbf{k}[0]})$$

$$11: \mathbf{v}' \leftarrow \mathbf{u} \cdot \mathbf{k}'$$

Householder Transformation of Encrypted Vectors.—Alg. 2 performs a key step in the Householder transformation, which reflects a vector about a given hyperplane, on an encrypted vector. For use in PCA, we need to choose a specific reflection hyperplane that transforms the input vector \mathbf{v} into a vector (of the same norm) with zeros in all coordinates except for the first. The output \mathbf{v}' of Alg. 2 represents this hyperplane; the Householder matrix obtained as $\mathbf{H} = \tilde{\mathbf{I}}^{(h \times h)} - 2\mathbf{v}' \times \mathbf{v}'^T$, where $\tilde{\mathbf{I}}$ is the identity matrix, satisfies that $\mathbf{H} \times \mathbf{v}$ has a nonzero element only in the first coordinate. This method is used in QR^T to iteratively apply orthogonal transformations to the input matrix to convert it into a lower triangular matrix. Following the standard technique, the norm of the input vector (computed in Lines 1–3) is added to or subtracted from its first coordinate (Line 7), depending on the sign of the first coordinate (Line 4) for numerical stability. Afterwards, the vector is normalized (Lines 9–11) to obtain the desired reflection vector.

Alg. 2 requires the evaluation of non-polynomial functions, including the sign function (alternatively, $g(x) = x/\sqrt{x^2}$, Line 4), the square root, and the inverse square root. To this end, SF-PCA applies Chebyshev polynomial approximation [89] to each function on a pre-determined input range (agreed upon in Step 1; §.5.2). In addition, we use the baby-step giant-step technique [90] to further reduce the complexity of evaluating degree- d polynomials, resulting in a multiplicative depth of $\lceil \log(d) + 1 \rceil$ and $2 \cdot \sqrt{2d} + \frac{1}{2} \log_2(d) + \mathcal{O}(1)$ ciphertext multiplications. We denote this quantity as $l(d)$ in our algorithms. We discuss the choice of approximation intervals in §.6.3. For the communication cost, we calculate the number of DBootstrap executions as the multiplicative depth of this method divided by the number of available ciphertext levels λ (§.3).

QR Factorization of Encrypted Matrices.—QR factorization decomposes an input matrix \mathbf{V} into an orthogonal matrix \mathbf{Q} and a lower-triangular matrix \mathbf{R} such that $\mathbf{V} = \mathbf{R} \times \mathbf{Q}$. This is repeatedly used in Steps 4, 6 and 7 of SF-PCA’s workflow. In Alg. 3, we describe both the transposed-QR factorization QR^T that is executed by one DP on an encrypted matrix and its distributed equivalent DQR^T . DQR^T performs a QR factorization in a federated manner on an encrypted matrix that is distributed among the DPs, requiring the DPs to aggregate (denoted by Ξ) their partial results in Lines 3, 5, and 18. In Step 4 of SF-PCA, QR^T is executed on a matrix with $h = m$ columns (i.e., same as the number of features), whereas DQR^T is executed on a matrix with $h = n$ columns distributed among the DPs, where each DP, has n_i columns. HH and the vector-matrix multiplications in Lines 5 and 18 are the only operations with a cost that depends on h . DQR^T requires more communication among the parties, and the complexity of QR factorization depends mainly on the number of rows δ , which is the same in both QR^T and DQR^T , and not on h . Hence, we use DQR^T only when the difference between n_i and m is large enough to compensate for the

communication overhead, i.e., when $\xi \cdot \log_2(n_{\max}) < \log_2(m)$, with a factor ξ determined by the properties of the network setup (e.g., latency). Note that $n_{\max} = \max_i(n_i)$.

Algorithm 3 - Encrypted \mathbf{QR}^T Factorization (or \mathbf{DQR}^T):

Input: Encrypted $\mathbf{V}^{(\delta \times h)}$

Output: Encrypted $\mathbf{Q}^{(\delta \times h)}$ and $\mathbf{R}^{(\delta \times \delta)}$, such that $\mathbf{R} \times \mathbf{Q} = \mathbf{V}$

Comp. Cost: $\mathcal{O}(\delta^2 + \delta \cdot \zeta_{\text{HH}})$ Mults, $\mathcal{O}(\delta^2 \cdot (1 + \log_2(h)) + \delta \cdot \zeta_{\text{HH}})$ Rots, where ζ_{HH} refers to the cost of HH (Alg. 2).

Comm. Cost: $C = \delta \cdot \zeta_{\text{HH}} + \frac{4\delta^2}{\lambda} \cdot \lceil \frac{h}{t} \rceil$ Cipher. (**DQR:** $C + 3 \cdot \delta \cdot \lceil \frac{h}{t} \rceil$ Cipher.)

<p>1: $\mathbf{H} \leftarrow \mathbf{0}^{(\delta \times h)}$</p> <p>2: for $i = 0, \dots, \delta - 1$ do</p> <p>3: $\mathbf{v} \leftarrow \text{HH}(\mathbf{V}[0, :])^T$ (DQR: Ξ in Line 2 of Alg. 2)</p> <p>4: $\mathbf{H}[i, :] \leftarrow \mathbf{v}^T$</p> <p>5: $\mathbf{v}' \leftarrow \mathbf{V} \times \mathbf{v}$ (DQR: $\Xi(\mathbf{v}')$)</p> <p>6: for $j = 0, \dots, \delta - i - 1$ do</p> <p>7: $\mathbf{V}[j, :] \leftarrow \mathbf{V}[j, :]$ $-2 \cdot (\mathbf{v}^T \cdot \text{Dup}_{(\delta-i)}(\mathbf{v}'[j]))$</p> <p>8: end for</p> <p>9: $\mathbf{r} \leftarrow \text{Rot}_{(-i)}(\mathbf{V}[0, :])$</p> <p>10: $\mathbf{R}[i, :] \leftarrow \mathbf{r}[:\delta]$</p> <p>11: for $j = 0, \dots, \delta - i - 1$ do</p>	<p>12: $\mathbf{V}[j, :] \leftarrow \text{Rot}_i(\mathbf{V}[j + 1, :])$</p> <p>13: end for</p> <p>14: end for</p> <p>15: $\mathbf{Q} \leftarrow [\tilde{\mathbf{I}}^{(\delta \times \delta)} \mathbf{0}^{(\delta \times (h - \delta))}]$</p> <p>16: for $i = \delta - 1, \dots, 0$ do</p> <p>17: $\mathbf{H}[i, :] \leftarrow \text{Rot}_{(-i)}(\mathbf{H}[i, :])$</p> <p>18: $\mathbf{h}' \leftarrow \mathbf{Q} \times \mathbf{H}[i, :]^T$ (DQR: $\Xi(\mathbf{h}')$)</p> <p>19: for $j = 0, \dots, \delta - 1$ do</p> <p>20: $\mathbf{Q}[j, :] \leftarrow \mathbf{Q}[j, :]$ $-2 \cdot \mathbf{H}[i, :]^T \cdot \text{Dup}_i(\mathbf{h}'[j])$</p> <p>21: end for</p> <p>22: end for</p>
---	---

From Lines 1 to 14, the input matrix \mathbf{V} is multiplied by the Householder matrix $\mathbf{H} = \tilde{\mathbf{I}} - 2\mathbf{v} \times \mathbf{v}^T$, where $\mathbf{v} = \text{HH}(\mathbf{V}[0, :])^T$ is the Householder vector obtained by Alg. 2 with the first row of \mathbf{V} as input. This transformation is recursively performed on the (i, i) minors of \mathbf{V} by discarding the first row and the first column to incrementally obtain the lower-triangular matrix \mathbf{R} . Due to SF-PCA's vectorized encoding scheme, the sub-matrix is efficiently obtained by applying a single ciphertext rotation per row (Line 9). In SF-PCA, \mathbf{R} is only used during the eigendecomposition in Step 6. \mathbf{Q} is computed in the second part (Lines 15 to 22) and corresponds to the product of all Householder matrices \mathbf{H} .

Recall that we minimize bootstrapping by refreshing only small-dimensional data that are globally shared among the DPs (§.5). The intermediate values in \mathbf{QR}^T satisfy this condition as they are derived from the input matrix that is already aggregated. Hence, the optimized number of invocations of $\text{DBootstrap}(\cdot)$ for \mathbf{QR}^T corresponds to its multiplicative depth divided by λ . For \mathbf{DQR}^T , the input matrix is split among the DPs. In this case, the results

of the collective aggregation (Ξ) in Lines 5 and 18, which constitute globally shared vectors among the DPs, are bootstrapped before being broadcast (shown as the additional cost).

Eigendecomposition of Encrypted Matrices.—Alg. 4 decomposes an encrypted matrix \mathbf{M} into $\mathbf{Q} \times \mathbf{L} \times \mathbf{Q}^T$, where \mathbf{Q} is a matrix of eigenvectors and \mathbf{L} is a diagonal matrix with the diagonal defined by the encrypted vector of eigenvalues l . The eigenvalues are ordered from the largest to the smallest. We adapt the standard QR iteration algorithm [79], [91] to the setting with an encrypted input matrix. The encrypted matrix is first tridiagonalized, i.e., transformed to a matrix where the only nonzero elements are in the diagonal, the subdiagonal, or the superdiagonal, which is known to improve the convergence rate of eigendecomposition [79]. The tridiagonalization is achieved by applying Householder transformations (using Alg. 2) to different subparts of the matrix to introduce zeros (Lines 2 to 11 in Alg. 4). The resulting encrypted matrix \mathbf{T} is then iteratively factorized using \mathbf{QR}^T (Line 17) into $\mathbf{R} \times \mathbf{Q}'$ (note the row-wise application of QR) and reconstructed as $\mathbf{Q}' \times \mathbf{R}$ to gradually transform the matrix into a diagonal matrix. During this process, the last diagonal element converges to the smallest eigenvalue of the input. This is then executed for each eigenvalue in an ascending order, and the corresponding eigenvectors are obtained from the product of all \mathbf{Q}' matrices. We perform all small-matrix multiplications (Lines 6, 7, 8, 18) by encoding each matrix in a single ciphertext and employing the technique of Jiang et al. [59]. We refer to this method as M5 to distinguish from the large-scale, unbalanced setting in M3 with ciphertext-cleartext multiplications. Multiplying two $s \times s$ encrypted matrices requires $5s$ Mults and $3s + 5\sqrt{s}$ Rots. We convert the matrices to our row-wise encoding scheme (in Lines 6, 9, 10, 12, 15, 17, 20, 22, 23) using one multiplication and one rotation per row, only to efficiently perform row and column selections. Similarly as Alg. 3, this method operates on globally shared inputs, and its optimized communication cost scales with the multiplicative depth divided by λ , in addition to the costs of the HH and \mathbf{QR}^T subroutines.

Algorithm 4 - Encrypted Eigendecomposition (Eigen):

Input: Encrypted symmetric $\mathbf{M}^{(\eta \times \eta)}$, number of iterations w

Output: Encrypted $\mathbf{Q}^{(\eta \times \eta)}$ and $l^{(1 \times \eta)}$, where the rows of \mathbf{Q} are eigenvectors of \mathbf{M} , and l has corresponding eigenvalues

Comp. Cost: $\mathcal{O}(\eta \cdot (1 + \zeta_{\text{HH}} + w \cdot \zeta_{\text{QR}}) + \eta \cdot (1 + w \cdot \zeta_{\text{M5}}))$ Mults and $\mathcal{O}(\eta \cdot (\zeta_{\text{HH}} + w \cdot \zeta_{\text{QR}}) + \eta \cdot (1 + (w \cdot \zeta_{\text{M5}})))$ Rots, where ζ_{HH} , ζ_{M5} and ζ_{QR} refer to the costs of HH, M5 and \mathbf{QR}^T .

Comm. Cost: $(\eta - 1) \cdot (\zeta_{\text{HH}} + w \cdot \zeta_{\text{QR}}) + \frac{(\eta - 1) \cdot (4 + 3w)}{\lambda} \cdot \left\lceil \frac{\eta}{t} \right\rceil$ Ciphertexts

1: $\mathbf{Q} \leftarrow \tilde{\mathbf{I}}^{(\eta \times \eta)}, \mathbf{T} \leftarrow \mathbf{0}^{(\eta \times \eta)}$ 2: for $i = 0, \dots, \eta - 3$ do 3: $\mathbf{v} \leftarrow \text{HH}(\mathbf{M}[0,1:i]^T)$ 4: $\mathbf{P} \leftarrow \tilde{\mathbf{I}}^{((\eta - i) \times (\eta - i))}$	13: for $i = \eta - 1, \dots, 1$ do 14: for $j = 0, \dots, w - 1$ do 15: $\mathbf{S} \leftarrow \mathbf{T}[i, i] \times \tilde{\mathbf{I}}[i, :]$ 16: $\mathbf{T} \leftarrow \mathbf{T} - \mathbf{S}$
--	--

<p>5: $P[1:1:] \leftarrow \tilde{T}^{((\eta-i-1) \times (\eta-i-1))} - 2 \cdot v \times v^T$</p> <p>6: $Q[i:, :] \leftarrow P \times Q[i:, :]$</p> <p>7: $PM \leftarrow P \times M$</p> <p>8: $M \leftarrow PM \times P^T$</p> <p>9: $T[i:i+2, i:i+2] \leftarrow M[:2, :2]$</p> <p>10: $M \leftarrow M[1:, 1:]$</p> <p>11: end for</p> <p>12: $T[\eta-2:, \eta-2:] \leftarrow M$</p>	<p>17: $Q', R \leftarrow \text{QR}^T(T)$</p> <p>18: $T \leftarrow Q' \times R$</p> <p>19: $T \leftarrow T + S$</p> <p>20: $Q[:i+1, :] \leftarrow Q[:i+1, :] \times Q'$</p> <p>21: end for</p> <p>22: $I[i] \leftarrow T[i, i]$</p> <p>23: $T \leftarrow T[:i, :i]$</p> <p>24: end for</p> <p>25: $I[0] \leftarrow T[0,0]$</p>
--	--

6.3. Non-Polynomial Functions on Encrypted Inputs

To approximate non-polynomial functions on chosen intervals, SF-PCA's default approach is to rely on homomorphic evaluations of Chebyshev polynomial approximations [90]. In Step 1 (Alg. 1), the DPs agree on the intervals and on the degree of the approximations. The complexity of the polynomial evaluation increases with the degree but is independent of the interval size, which influences the precision. While any interval selection approach may be used with SF-PCA, the approach we adopt in our evaluation in §.7 is for a DP (e.g., the one coordinating the collaboration or the one with the highest number of local samples) to set the intervals based on the estimated range of intermediate values to be encountered by running RPCA on a simulated dataset, obtained by upsampling its local data to match the size of the joint data. In §.8, we discuss an extension to SF-PCA that enables it to switch to secret sharing for the evaluation of non-polynomial functions, for which efficient bit-wise protocols exist for scaling the input to a common range for approximation. This effectively removes the need to choose intervals and, depending on the parameters, can further improve SF-PCA's accuracy (Appendix D.1).

7. System Evaluation

We show that SF-PCA, enabled by our optimization techniques (§.5), efficiently computes a PCA on high-dimensional inputs distributed among a large number of DPs. We demonstrate SF-PCA's practicality and accuracy on various datasets with the number of features ranging from 8 to 23,724 and including up to 60,000 samples. SF-PCA consistently obtains PCs that are highly similar ($r^2 > 0.9$) to those obtained by a standard non-secure PCA. SF-PCA also outperforms alternative privacy-preserving approaches in terms of accuracy and runtime, and offers stronger security guarantees compared to some. In §.7.7, we show that, contrarily to meta-analysis, SF-PCA remains accurate regardless of potential differences in the data distribution among the DPs.

7.1. Formal Analysis of Costs

SF-PCA's communication cost depends mainly on the number of features m , the number of components ρ and the number of power iterations p . SF-PCA's computation cost depends on the same parameters and optionally on the number of samples per DP n_i . For both, the

overall cost is amortized over the ciphertexts due to packing and the SIMD property of HE, effectively dividing the contributions of m and n_i to the complexity by the ciphertext capacity t . In Tab. 1, we show the theoretical costs for a single DP (DP_i) for each step in SF-PCA (Alg. 1).

The communication in Step 1 is due to the generation of the public key pk and evaluation keys $evks$ (including a relinearization key and $\log_2(t)$ rotation keys). All rotations in SF-PCA are executed by combining rotations of power-of-two shifts using the pre-generated keys. DBootstrap requires each DP to transmit and receive the equivalent of a ciphertext, and to perform one ciphertext addition (at a negligible cost). In the remaining steps, we analyze the communication cost in terms of the number of DBootstrap invocations, which depends on the cryptographic parameters and the number of multiplications to perform in each routine (§.6.2). In turn, the number of multiplications depends on the input dimensions, the degree of polynomial approximations and, for Eigen, the number of iterations w .

SF-PCA's overall communication cost is independent of the number of samples and is dominated by the bootstrapping execution. We optimize the performance of SF-PCA by selecting the computation approach with the lowest complexity, e.g., by choosing *Precomp* (whose complexity is independent of n_i) if the number of samples is large.

7.2. Implementation Details and Evaluation Settings

We implemented SF-PCA in Go [92], building upon Lattigo [93] and Onet [94], which are open-source Go libraries for lattice-based cryptography and decentralized system development, respectively. The communication between DPs is through secure TCP channels (using TLS). We evaluate our prototype based on a realistic network emulated using Mininet [95], with a bandwidth of 1 Gbps and a communication delay of 20ms between every two nodes. Unless otherwise stated, we uniformly and horizontally distribute the input data among 6 DPs. We deploy each DP on a separate Linux machine with Intel Xeon E5-2680 v3 CPUs running at 2.5 GHz with 24 threads on 12 cores and 256 GB of RAM. We provide the default system parameters of SF-PCA considered in our evaluation in Appendix B.

7.3. Microbenchmarks for MHE Protocols in SF-PCA

In Fig. 2, we summarize the runtimes for SF-PCA's main ciphertext operations as well as high-level linear algebra routines. Recall that each ciphertext contains up to $t = 2^{13}$ values and any operation is concurrently executed on all encrypted values. Multiplying a cleartext with a ciphertext is almost 8x faster than multiplying two ciphertexts with the default parameters. The transmission time of a ciphertext ($\text{Send}(c)$) depends mostly on the communication delay (20ms in our setting). In our default setting, DKeyGen takes 9 seconds to generate the public key, relinearization key, and 13 rotations keys.

7.4. Practical Scalability of SF-PCA Performance

We evaluated SF-PCA's scalability on simulated datasets of varying sizes. In Fig. 3, we show that SF-PCA's runtime (when computing eight PCs with ten power iterations) remains

almost constant when the dimensions are smaller than the ciphertext capacity t (set to 8,192 by default). It then grows linearly with the number of ciphertexts, i.e., with the number of features (m) and samples per DP (n_i) divided by t . Note that, since the protocol is synced among the DPs at each aggregation step, SF-PCA's runtime depends on the slowest DP, e.g., the DP with the largest local dataset, as shown in Fig. 3 and Fig. 7c in appendix. In all figures, we omit the negligible execution times of Steps 1 to 3. These steps require mostly non-iterative cleartext operations. In Fig. 3 (left panel), we set $n_i = 1,024$ and show that all SF-PCA's approaches (i.e., *Precomp* and *Seq* with QR^T or DQR^T) similarly scale linearly with m . *Precomp* is the most efficient approach for this range of values for m and $n_i = 6,144$ but becomes impractical with a large m . In these experiments, we found DQR^T to be consistently inferior to QR^T as the former's communication overhead overshadows its computational speedup. This is expected, since the computational gain of using DQR^T depends on how much smaller n_i is with respect to m (see §.6.2). This difference is never large enough to compensate for the communication overhead in our settings. The results in Fig. 3 (right panel), for which m is set to 256, show that SF-PCA's runtime remains constant when using *Precomp*, which does not depend on the number of samples.

We remark that SF-PCA's runtime is dominated by the time dedicated to the communication between the DPs (Fig. 4). The communication overhead ranges between 90% of the runtime with small input dimensions, i.e., when the packing capacity of the cryptoscheme is less exploited, and 45% of the runtime when the dimensions are equal or larger than t . Although SF-PCA is able to minimize its computation runtime with optimized federated and parallelized computation methods, its communication overhead is bounded by the available communication network. When the number of DPs (s) doubles, SF-PCA's runtime increases only by a factor of around 1.1. This is because the amount of local computation does not grow with s (Table 1) and because the cost of interactive routines only slightly increases with s . Based on Fig. 4, we estimate practical runtimes even for hundreds of DPs, e.g., 110 minutes for 200 DPs with a maximum of 1,024 data samples per DP. We discuss in §.8 how SF-PCA can be extended to handle availability issues given many DPs. In Fig. 4, SF-PCA's runtime grows linearly with the number of components in all its steps except in Step 6, where the eigendecomposition cost depends on the small matrix dimensions: $\rho \times \rho$. SF-PCA's runtime increases linearly with the number of power iterations; however, this parameter typically does not grow with the data size for RPCA.

In our default scenario, SF-PCA's runtime is multiplied by a small factor of 1.1x when the available bandwidth is halved and the communication delay doubled. Moreover, each ciphertext accounts for 2.5 MB, thus executing SF-PCA on a dataset with 8,192 features (or less) requires each DP to send 3.8 GB, independently of the number of samples n (which can be large).

7.5. Accuracy of SF-PCA Results

We demonstrate SF-PCA's accuracy and practicality on six real datasets, including MNIST [54] and two genomic datasets [55], [96] with thousands of patients and up to 23,724 features (see Appendix E for dataset details). We evenly and randomly split each dataset

among the DPs. In §.7.7, we show that SF-PCA computes the same results regardless of the data distribution among the DPs. In Fig. 5, we show that SF-PCA and the cleartext non-secure centralized Randomized PCA (RPCA, Fig. 1) achieve similar accuracy (according to the mean-squared error, MSE; and Pearson Correlation Coefficient, r^2), with respect to the PCs obtained using the standard non-secure PCA, i.e., the RPCA implementation provided by the *sklearn* Python package [97].

7.6. Comparison with Existing Works

We compare SF-PCA with existing approaches for federated or multiparty PCA, which we categorize into meta-analysis, centralized HE (C-HE), and secret sharing-based SMC solutions. A more detailed review of these approaches is provided in §.2.

Meta-analysis.—For comparison, we replicate the meta-analysis approach of Liang et al [41], whereby a central computing server performs a truncated SVD on the combined SVD results obtained independently by each DP. In Fig. 5, we show that this solution yields the least accurate results across all datasets. Note that SF-PCA significantly improves upon the accuracy of meta-analysis by emulating a centralized PCA. Moreover, most meta-analysis solutions [33], [34], [35], [36], [37], [38], [39], [40], [41], [42], [43] are not end-to-end secure as the DPs' intermediate results are revealed to an aggregator server (or to the other DPs). These solutions achieve similar runtimes as non-secure centralized solutions because they also operate on unprotected cleartext data.

Centralized HE (C-HE).—We estimate the runtime of an HE-based centralized solution based on SF-PCA's runtime as follows. We account for the fact that the computations cannot be distributed among the DPs and that all operations must be performed on the encrypted data. Recall that SF-PCA exploits local cleartext operations to optimize computation (e.g., §.6.1) and that multiplying two ciphertexts is 8 times slower than a plaintext-ciphertext multiplication. We also include the overhead brought by a centralized bootstrapping routine [90], which is two orders of magnitude slower than DBootstrap, e.g., 26 seconds for [90] vs. 0.49 seconds with DBootstrap. Furthermore, since centralized bootstrapping consumes levels and lowers the number of available levels for multiplications, C-HE would require more conservative cryptographic parameters with larger ciphertexts, and thus higher computation and communication costs. In Fig. 5, we show the estimated lower bound of the runtime for a C-HE solution executed by a single DP. We remark that SF-PCA, by distributing its workload and relying on efficient interactive protocols, is consistently 1–2 orders of magnitude faster than a C-HE solution. We note that we consider C-HE solutions based on the same underlying scheme of SF-PCA with comparable parameters, and that more sophisticated centralized solutions could be devised. However, those would still suffer from a high communication overhead and introduce a single point of failure due to the data centralization.

Secret Sharing-based SMC.—In Fig. 5, we compare SF-PCA's runtime with the linear (additive) secret sharing-based SMC solution proposed by Cho et al. [45]. In this solution, two computing servers perform PCA on secret-shared data, and a third server is responsible for the generation and distribution of correlated random numbers used in SMC protocols

(e.g., Beaver triples [101]). This additional party is trusted to correctly generate these values and not to collude with any other party. We ran Cho et al.'s publicly available, two-party solution [102] in our evaluation environment. We further estimated the runtime of this solution with 6 DPs under linear scaling with the number of DPs. We observe that SF-PCA is between 3x and 10x faster than the SMC solution while operating in a stronger threat model without the need for an honest third party. We also note that the SMC solution requires the entire dataset to be secret-shared among the computing parties, which can be costly for large datasets and complicate regulatory compliance. For example, with the Lung dataset [55], this represents a communication overhead of more than 60 GB. Finally, we note that SMC solutions heavily rely on interactive computations, leading to many rounds of communication in total. Since a large portion of SF-PCA is local non-interactive computation by each DP, SF-PCA remains practical even in constrained networks with high communication delays, unlike the SMC solutions. For example, when we double the delay from 20ms to 40ms, we observed that SF-PCA's runtime remains almost constant, whereas the SMC solution becomes 1.9 times slower in the two-party setting. In §.8, we describe an extension of SF-PCA which uses secret sharing specifically for non-polynomial operations over low-dimensional inputs.

7.7. Example Application of SF-PCA in Genomics

To further demonstrate the utility of SF-PCA, we used it to analyze a genomic dataset of 2,504 individuals with 1,773 features (a subset of genetic variants from chromosome 20). PCA is a standard step in many genomic analysis workflows, e.g., in genome-wide association studies [9], for capturing ancestry patterns in a dataset. We split the data among three DPs such that each DP only has samples belonging to a specific ancestry group (Fig. 6.d). The plots show individual samples projected onto the first two PCs. Consistent with the quantitative evaluation in §.7.5, SF-PCA (Fig. 6.b) is able to accurately identify the low-dimensional structure spanned by the data samples, almost exactly replicating the output of a centralized cleartext PCA on the full dataset, independently of how the data is split among the parties (Fig. 6.a). The meta-analysis approach for PCA (Fig. 6.c) results in a distorted data landscape due to the limited view of each DP. In Fig. 6.d, we highlight the output of SF-PCA that is visible to one of the DPs; while all DPs obtain projected data according to a *unified* subspace identified by the PCA, each DP sees only a portion of the output associated with their local data as required by our security model.

8. Extensions

SF-PCA can be extended in several ways to incorporate additional features. First, SF-PCA's multiparty construction enables it to seamlessly and securely (i.e., without decryption) switch between MHE and secret sharing-based SMC [30], [31], [103] (see Appendix D.1). This enables SF-PCA to leverage more efficient and accurate protocols to evaluate non-polynomial functions (e.g., sign tests) on small-dimensional inputs, while using MHE for operations over large encrypted vectors and matrices where the SIMD property of MHE leads to efficient performance with minimal communication. Next, the modular design of SF-PCA enables its federated routines to be used to perform RPCA on vertically partitioned data (Appendix D.2). SF-PCA could also be extended to provide differential

privacy (Appendix D.3), although setting a meaningful privacy parameter may be difficult, by incorporating an interactive protocol in which the DPs sequentially shuffle an encrypted list of noise values before adding them to the results upon decryption [104]. Lastly, to cope with the possibility of a subset of DPs becoming unavailable during the PCA computation—particularly relevant for the setting with many DPs, SF-PCA can be instantiated with a threshold secret sharing of the MHE secret key [53] to allow a subset of DPs to continue the protocol execution (Appendix D.4).

9. Discussion and Conclusions

We introduced SF-PCA, a decentralized system for securely and efficiently executing PCA on data held by multiple data providers. SF-PCA ensures input confidentiality as long as at least one DP is honest. Furthermore, the local private data never leave the DPs’ premises given the federated design of SF-PCA. Our system builds on a range of optimized MHE-based routines we developed for key computational operations in PCA such as large-scale cleartext-ciphertext matrix multiplications and sophisticated linear algebra transformations, including matrix factorization and orthogonalization. SF-PCA obtains accurate results within practical runtimes on large matrices including tens of thousands of features, and efficiently scales with the number of data providers and the input dimensions due to our optimization strategies.

Our work shows that an end-to-end secure solution for high-complexity data analysis tasks such as PCA is practically feasible. Incorporating SF-PCA into existing privacy-preserving federated analysis methods (e.g., see Appendix G) and deploying it in a range of practical applications are natural next steps for our work. Our design principles and optimization techniques that have led to the practical performance of SF-PCA, as well as the optimized MHE routines for key linear algebra operations such as eigendecomposition are broadly applicable to other problems in federated analytics.

Acknowledgements

We thank Louis Vialar and the reviewers for their comments. This work was partially supported by NIH R01 HG010959 (to B.B.) and by NIH DP5 OD029574, RM1 HG011558, and Broad Institute’s Schmidt Fellowship (to H.C.). J.R.T.-P. and J.-P.H. are co-founders of the start-up Tune Insight. All authors declare no other competing interests.

Appendix A: CKKS

We instantiate SF-PCA’s multiparty scheme with the Cheon-Kim-Kim-Song (CKKS) cryptosystem [61]. CKKS parameters are denoted by the tuple $(\mathcal{N}, \Delta, \eta, mc)$, where \mathcal{N} is the ring dimension; Δ is the plaintext scale by which any value is multiplied before it is quantized and encrypted/encoded; η is the standard deviation of the noise distribution; and mc represents a chain of moduli $\{q_0, \dots, q_L\}$ such that $\prod_{i \in \{0, \dots, \kappa\}} q_i = Q_\kappa$ is the ciphertext modulus at level κ , with Q_L the modulus of fresh ciphertexts. Operations on a ciphertext c at level κ and scale Δ with $\Delta < Q_\kappa$ are performed modulo Q_κ . We denote by $\{c, L, \Delta\}$, with $c = (c_0, c_1) \in R_{Q_L}^2$, and $\tilde{p} \in R_{Q_L}$, a fresh ciphertext at level L with scale Δ and a plaintext, respectively.

Appendix B: Symbols & Default Values

TABLE 3:

Glossary of Symbols and Their Default Values in SF-PCA.

Symbol	Definition	Default
s, p, w	# DPs, # power and eigen iterations	6, 10, 5
$\psi + \alpha = \rho$	# PCs + oversampling = # components	4 + 4 = 8
ζ^*, ζ, d	Optimized cost, cost, approx. degree	-, -, 31
m, n, n_i	# features, # samples tot. & at DP_i	$2^8, 6144, 2^{10}$
\mathcal{N}, λ	Ring dim., # available levels	$2^{14}, 7$
R_Q	Plain/Ciphertext domain	-
c	encrypted vector/ fresh ciphertext with $c = (c_0, c_1) \in R_{Q_L}^2$	-
$\tilde{p} \in R_{Q_L}, sk, pk$	plaintext, secret & public keys	-, -, -
t, \bullet	c capacity, dot product	$2^{13}, -$
$\mathbf{M}^{(a \times b)}, \overline{\mathbf{N}}^{(b \times c)}$	Generic encrypted and cleartext matrices	-, -
$\mathbf{M}[i, j], v[i]$	Matrix/vector elem. at index $(i, j)/i$	-, -

Appendix C: Security Analysis

We rely on the real/ideal simulation paradigm [105] to show that SF-PCA achieves the input confidentiality requirement defined in §.4. A computationally bounded adversary that controls up to all but one DP cannot distinguish a *real* world experiment, in which the adversary is given actual data from an execution of our protocol from the views of the compromised DP(s), and an *ideal* world experiment, in which the adversary is given random data generated by a simulator.

The semantic security of the CKKS scheme used in SF-PCA is based on the hardness of the decisional-RLWE problem [61], [80], [106]. Mouchet et al. [48] proved that their distributed protocols, i.e., DKeyGen and DKeySwitch, are secure under the simulator paradigm. They show that the distribution of the cryptoscheme preserves its security in the passive-adversary model with all-but-one dishonest DPs, as long as the decisional-RLWE problem is hard. Their proofs are based on the BFV cryptoscheme; Froelicher et al. [24] showed that the proofs still hold with CKKS, as the same computational assumptions hold, and the security of CKKS is based on the same hard problem as BFV. They make a similar argument for DBootstrap and prove its security. The security of the cryptoscheme used by SF-PCA follows from these results.

Proposition 1.

Assume that SF-PCA uses CKKS encryptions with parameters $(\mathcal{N}, \Delta, \eta, mc)$ ensuring post-quantum security. Given a passive adversary corrupting at most $s - 1$ parties out of s parties in total, SF-PCA achieves input confidentiality.

Sketch of the Proof.

We consider a real-world simulator \mathcal{S} that simulates the view of a computationally-bounded adversary corrupting $s - 1$ parties, i.e., it has access to the inputs and outputs of $s - 1$ parties. In Step 1 of SF-PCA's workflow (Alg. 1), the simulator obtains the public parameters and the entire matrix $\tilde{\mathbf{A}}$, except the rows that belong to the honest DP. From Step 1 to the end, the DPs exchange only collectively encrypted information. In Step 8, each DP projects its local data on the obtained collectively encrypted PCs. If required by the application, the collectively encrypted result is switched to each DP's public key so that they can decrypt the final result. To avoid information leakage about data and/or about the encryption keys from the processed ciphertexts [84], we rely on existing countermeasures [48], [85], [86] and add fresh noise (i.e., rerandomizing) sampled from a distribution that has a variance significantly larger than that of the input ciphertext's noise distribution to the processed ciphertext [84]. Alternatively, this result can be kept encrypted and used for future steps without ever being decrypted. Hence, by generating random ciphertexts with parameters $(\mathcal{N}, \Delta, \eta, mc)$, \mathcal{S} can simulate all the values communicated during the entire process such that the real outputs cannot be distinguished from the ideal ones. The sequential composition of all cryptographic functions remains simulatable by \mathcal{S} as there is no dependency between the random values that an adversary can exploit. Also, the adversary cannot decrypt collectively encrypted data unless all DPs collude, which would contradict the considered threat model (§.4). Following this, SF-PCA ensures *input confidentiality* for the honest DP(s).

Appendix D: Extensions

D.1. Hybrid Use of Multiparty Security Primitives

SF-PCA's multiparty construction enables it to seamlessly switch between MHE and secure multiparty computation (SMC) primitives based on secret sharing [30], [31], [103]. Intuitively, an MHE ciphertext is transformed to linear (additive) secret shares (LSS) through a collective masked decryption by the DPs, i.e., each DP partially decrypts the ciphertext and masks the result with its secret share, whereas the last DP decrypts and obtains its share. After the computations in the LSS domain, to transform the result back to an MHE ciphertext, each DP encrypts its local share of the result such that it can be aggregated under MHE with all DPs' encrypted shares. We detail these procedures in Protocol 1. SF-PCA always employs MHE to execute large-dimensional matrix operations and can perform non-polynomial operations, e.g., square root and divisions (in Alg. 2), as well as small-matrix operations (in Alg. 4), using LSS-based routines [45]. This combines the strengths of both approaches: On the one hand, relying on edge-computing and the SIMD property of MHE, SF-PCA efficiently performs vectorized and parallel operations over large encrypted matrices while minimizing communication. On the other hand, relying on LSS-based SMC, SF-PCA simplifies its usage by removing the need to choose intervals for non-polynomial function approximations. Note that efficient protocols exist for computing the bit-length of a secret-shared value [107], which can be used to map the input to a common interval for accurate approximation. In addition, LSS-based routines can be more efficient for computation over small data, e.g., eigendecomposition of a tiny matrix in RPCA, where the ciphertext packing is underutilized for MHE. In SF-PCA, all costly non-polynomial

operations (e.g., see Alg. 2) are executed on a single scalar input. Thus, executing these operations on compact secret-shared data can further reduce the computational cost of SF-PCA.

LSS Scheme.

We implemented a collection of SMC routines used by the prior work on LSS-based PCA [45] to perform the required non-linear operations (comparison, square root, and division), newly extending the support to more than two DPs and 128-bit security. These protocols build upon a combination of existing SMC techniques [31], [101], [103], [107], [108], [109], [110]. All values and shares are encoded as field elements \bar{x} (or $\bar{\mathbf{x}}$ for a vector of elements) in $\mathbb{Z}_{\bar{p}}$, with \bar{p} a prime, by relying on a fixed-point representation [108]. For example, with 2 parties, $x \in \mathbb{Z}_{\bar{p}}$ is secret shared as $r \in \mathbb{Z}_{\bar{p}}$ and $(x - r) \in \mathbb{Z}_{\bar{p}}$. Additions consist in simple share additions, whereas multiplications are done by relying on Beaver multiplication triples [101]. Following the prior work, we adopt the server-aided model of preprocessing whereby a third-party generates these triples to facilitate the main interactive computation with efficiency. This scheme can be modified to avoid the need of a trusted node at setup, thus following SF-PCA's default threat model, by relying on an interactive protocol for the setup to be executed among all DPs. Adapting existing solutions [30], [48] to SF-PCA is part of future work.

Protocol to Switch Between MHE and LSS.

We build on the collective bootstrapping protocol [48]. We split this protocol in two rounds and add a conversion to/from the field $\mathbb{Z}_{\bar{p}}$ of the LSS scheme, see Protocol 1. We assume that DP_1 wants to perform a function f_{LSS} on the encrypted vector c . The security of the protocol can be derived from the security of the original DBootstrap (for which Froelicher et al. [24] prove that statistical indistinguishability is preserved as long as the masks are sampled from the correct distribution), from the LSS scheme guarantees (i.e., statistical indistinguishability), and from the security of CKKS and the hardness of the decisional-RLWE problem [61], [80], [106].

Evaluation.

Switching to LSS removes the need for defining approximation intervals to evaluate non-polynomial functions hence simplifies the usage of SF-PCA. Depending on the setting, it can also improve SF-PCA's accuracy. The intervals for the non-polynomial operations depend on the DPs' data and, as SF-PCA's intermediate results are repeatedly orthogonalized (through QR^T), these ranges can be accurately inferred upfront by the DPs, e.g., by simulating the protocol (§.6.3).

Protocol 1

MHE \iff LSS

Input: DP_1 has $\mathbf{c}_{pk} = (c_0, c_1) = \{c, \tau, \Delta\} \in R_{Q_r}^2$ a ciphertext encrypting $\tilde{\mathbf{p}}$. ν is a security parameter, sk_i the secret-key of each DP_i , χ_{err} a distribution over R , where each coefficient is independently sampled from Gaussian distribution with the standard deviation $\sigma = 3.2$, and bound $\lfloor 6\sigma \rfloor$. $\text{Encode}(\cdot)$ is the mapping from a plaintext

encoded in R to the equivalent encoding in $\mathbb{Z}_{\bar{p}}$. Let T be the bound on all possible coefficients in the polynomial representation of \bar{p} encoding real data values and l be the bound on the possible bit length of real data values encoded in $\mathbb{Z}_{\bar{p}}$.

Output: $c'_{pk} = \{c', L, \Delta\}$

Constraints: $Q_\tau > (s+1) \cdot T \cdot 2^v \ \& \ (s+1) \cdot 2^{v+l} < \bar{p}$

- 1: DP_1 broadcasts $\{c, \tau, \Delta\}$
- 2: Each DP_i for $i = 2, \dots, s$
- 3: Samples $\mathbf{a}_i \leftarrow \text{Uniform}(R_{T \cdot 2^v})$, $e_i \leftarrow \mathcal{X}_{err}$
- 4: Sends $\mathbf{h}_i = sk_i \cdot c_1 + \mathbf{a}_i + e_i \text{ mod}(Q_\tau)$ to DP_1
- 5: Assigns $\bar{\mathbf{a}}_i = \text{Encode}(-\mathbf{a}_i) \text{ mod}(\bar{p})$
- 6: DP_1 :
- 7: Samples $\mathbf{a}_1 \leftarrow \text{Uniform}(R_{T \cdot 2^v})$, $e_1 \leftarrow \mathcal{X}_{err}$
- 8: Computes $\mathbf{h}_1 = sk_1 \cdot c_1 + \mathbf{a}_1 + e_1 \text{ mod}(Q_\tau)$
- 9: Computes $\mathbf{h}' = c_0 + \sum_{i=1}^s \mathbf{h}_i \text{ mod}(Q_\tau)$
- 10: Assigns $\bar{\mathbf{a}}_1 = \text{Encode}(\mathbf{h}' - \mathbf{a}_1) \text{ mod}(\bar{p})$
- 11: All DP_i :
- 12: Compute $\bar{\mathbf{r}}_i = f_{LSS}(\bar{\mathbf{a}}_i)$
- 13: Compute $\bar{\mathbf{r}}_i = \bar{\mathbf{r}}_i - \bar{\mathbf{b}}_i \text{ mod}(\bar{p})$, with $\bar{\mathbf{b}}_i \leftarrow \text{Uniform}(\mathbb{Z}_{2^{v+l}})$
- 14: Encrypt $c'_{pk} = \text{Enc}(pk, \bar{\mathbf{b}}_i)$
- 15: Each DP_i for $i = 2, \dots, s$: Sends $c'_{pk}^{(i)}$ and $\bar{\mathbf{r}}_i$ to DP_1
- 16: DP_1 :
- 17: Computes $\bar{\mathbf{r}}' = \sum_{i=1}^s \bar{\mathbf{r}}_i \text{ mod}(\bar{p})$ and encrypts $c'_{pk} = \text{Enc}(pk, \bar{\mathbf{r}}')$
- 18: Computes $c'_{pk} = c'_{pk} + \sum_{i=1}^s c'_{pk}^{(i)}$

For example, with the MNIST dataset and the parameters of Fig. 5, the DPs define 16 distinct intervals to evaluate 1,047 polynomial approximations. This is because the ranges of values are constant across the dimensions and across the iterations of the same operations. For the same dataset, relying on SF-PCA +LSS improves the Pearson correlation between SF-PCA's PCs and the PCs obtained with a standard non-secure centralized PCA from 0.91 to 0.92 (when using our default parameters, Tab. 3). SF-PCA's accuracy depends on the size and degree of the intervals hence can be improved by refining these parameters. We illustrate this on the execution of a single QR^T in Fig. 7b. We note that QR^T is used (iteratively) in Steps 4, 6 and 7 of SF-PCA and that all non-polynomial operations in SF-PCA are executed in the Householder (HH, Alg. 2) that is called in line 2 of QR^T . For QR^T on a 8×8 matrix, HH is called seven times and requires the evaluation of three non-polynomial functions. In SF-PCA, this requires the definition of 21 approximation intervals, i.e., one per non-polynomial function. We show that using a single large interval (with a polynomial of degree 63; [0:1000;63]) for all operations already yields results that are correlated with the results obtained by a cleartext solution. SF-PCA's accuracy can then be improved by either downsizing the interval (to [0:100;63]), increasing the approximation

degree (to $[0:1000;127]$), or by using more fine-grained intervals for the different steps in the computation ($[0:100;63]$ for the first execution of HH and $[0:1;63]$ afterwards).

In Fig. 7a, we show that SF-PCA's runtime is similar with or without this extension. SF-PCA's computational cost is reduced by computing on secret shares, instead of on encrypted vectors, but this gain is overshadowed by the communication overhead brought by both the protocol for switching between MHE and LSS and by LSS distributed computations. SF-PCA scales similarly with its default approach (SF-PCA in Fig. 7a) and when switching to LSS for non-polynomial and small-dimensional operations (SF-PCA +LSS). Switching to LSS only for non-polynomial operations (SF-PCA +LSS-OP) is around 1.4x slower than SF-PCA +LSS due to the communication overhead brought by the high-number of switches between the two schemes. SF-PCA can optimize its runtime for the small-dimensional eigendecomposition (Step 6) by performing it entirely in the LSS domain, which is up to 1.5x faster than in its basic approach in this scenario. When operating on larger dimensions, i.e., in Step 4 (Alg. 1), SF-PCA only switches to LSS for small-dimensional (i.e., single value as shown in Alg. 2) non-polynomial operations as this can improve its precision. Performing sequences of operations in LSS in Step 4 would require to switch and operate on large-dimensional secret-shared elements, which would further increase the communication overhead. In Fig. 7b, we show that the runtimes of most LSS operations are in the same order of magnitude as MHE operations.

D.2. Vertically Partitioned Data

SF-PCA's workflow can be easily adapted to work with a vertically partitioned input matrix by modifying the interactive computation among the DPs while leveraging the same local operations as before. Because of the different way the data is split, some of the DPs' intermediate results have to be combined (aggregated or concatenated) at different points in SF-PCA's workflow; this does not change the nature of the underlying operations that are optimized in the default setting of SF-PCA. In the vertical case, the overall computation and communication complexities depend on the total number of samples n and the number of features per DP m_i , whereas these depend on n_i and m in SF-PCA's original approach. We show both approaches in Fig. 8. The mean-centralization is up to eight times less expensive than in SF-PCA's original workflow, because each DP_{*i*} keeps its part of the averages' vector \tilde{o}_i in cleartext.

D.3. Differential Privacy

SF-PCA can be extended to provide differential privacy by leveraging an interactive protocol in which the DPs sequentially shuffle an encrypted list of noise values before adding them to the results upon decryption [104]. The choice of privacy parameters and maintaining accuracy are part of future work.

D.4. Fault Tolerance

To cope with the possibility of a subset of DPs becoming unavailable during the PCA computation, which is particularly relevant when there are many DPs, SF-PCA can be

extended by employing a \hat{s} -out-of- s threshold secret-sharing for the MHE secret keys [53], where s is the number of DPs. Note that the main setting of SF-PCA considers $\hat{s} = s$. Setting \hat{s} to be smaller than s changes SF-PCA’s security model to tolerate up to $\hat{s} - 1$ dishonest DPs. As long as at least \hat{s} DPs are available for each interactive step, SF-PCA’s execution continues without interruption. In certain steps of SF-PCA the omission of a subset of parties may result in their local data not being accounted for in the computation. However, given the iterative nature of the RPCA algorithm, the overall results are expected to be robust against such omissions with a sufficient number of iterative steps.

Appendix E: Datasets

The Wine dataset [98] contains 4,898 wine samples with physicochemical attributes as features and a quality score as label. The Lung dataset [55] contains 9,098 patients with 23,724 genomic variations (as features) and a label indicating the presence of a cancer. The PIMA dataset (768×8) [99] contains medical observations collected from an Indian community that can be used to predict the presence of diabetes. Chr20 (2,502×1,773) [96] is a subset of the genomic data available in the 1,000 Genomes dataset. In the MNIST dataset (70,000×784) [54], each sample describes the grey-scale image of a single handwritten digit. Finally, the Vehicle [100] dataset contains 19 features extracted from each of the 435 images of buses or cars.




Figure 7:

In Fig. 7a, we show SF-PCA’s runtime when using LSS extension. In Fig. 7b, [0 : 100;63] indicates that the approximations are done in an interval between 0 and 100 with a degree of 63. Fig. 7c depicts SF-PCA’s runtime when one DP has 32,786 data samples and the remaining samples are evenly split among 5 DPs.

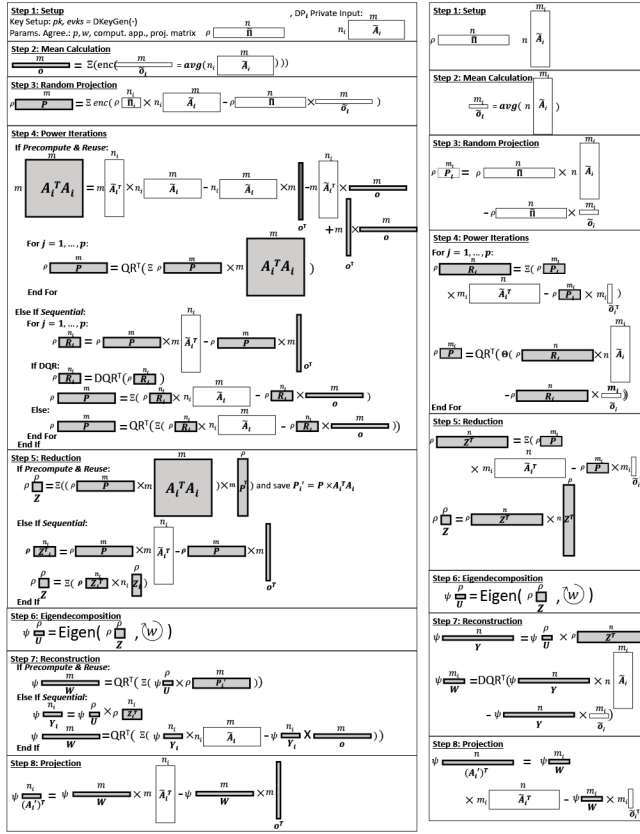


Figure 8: SF-PCA’s secure workflow with horizontally (left) and vertically (right) partitioned data.

The execution is depicted from DP_i’s point of view. The filled boxes indicate encrypted matrices and the empty boxes show cleartext matrices. The dimensions are shown with the box sizes and are indicated on the left and top of the corresponding box. Ξ indicates a collective aggregation, Θ a collective concatenation, and, in both cases, the result is broadcast to all DPs. The dimensions equal to 1 are omitted and vectors replicated to comply with the matrix-multiplication dimensions are not shown. A tilde indicates cleartext.

Appendix F: Runtime Scales with the Slowest DP

We show in Figure 7c that SF-PCA’s runtime depends on the maximum number of local samples among the DPs. In this example, the DP with the maximum number of samples has 32,768 samples and the other samples are evenly split among the remaining 5 DPs. Even as the total number of data samples increases, SF-PCA’s runtime remains constant since the maximum number of local samples stays the same.

Appendix G: Using SF-PCA to Improve Machine Learning Efficiency and Accuracy

By combining SF-PCA with a privacy-preserving solution for a downstream machine learning (ML) task, a secure federated ML workflow supporting the full analytic pipeline, encompassing pre-processing (e.g., dimension reduction), training, and inference, can

be built. For example, SF-PCA can be seamlessly integrated with existing MHE-based solutions for training generalized linear models [24], [27] or neural networks (NNs) [26]. As the training time of these solutions increase with the number of features in the dataset, SF-PCA may be a useful solution for reducing the scale of high-dimensional datasets to speedup model training. For example, executing SF-PCA on the MNIST dataset (see Fig. 5) to project it on 5 PCs takes 1 hour and reduces the number of features by a factor of 152. Training a model using the PCs instead of the original features would reduce by a factor 7 the runtimes of previously mentioned secure solutions. Such an approach can also improve the accuracy of ML models when dimension reduction results in noise removal and more informative features, especially in limited data settings [111], [11], [12], [13]. We illustrate this use case by training a NN model (i.e., a multilayer perceptron with hidden layer made of four nodes, sigmoid activation functions, and one output node) to perform classification on the Vehicle dataset [100], which contains 435 samples with 18 features derived from vehicle images (Appendix E). We observed that the model trained without any preprocessing achieves a prediction accuracy of 55% on the test set, whereas training the model on 5 PCs obtained by SF-PCA (applied to the training data) as features yields an accuracy of 87%, which increases to 97% with 8 PCs. This small example illustrates the fact that by de-correlating the features and reducing their number, PCA can improve ML model accuracy.

References

- [1]. Hotelling H, "Analysis of a Complex of Statistical Variables into Principal Components." Journal of educational psychology, 1933.
- [2]. Pearson K, "LIII. On Lines and Planes of Closest Fit to Systems of Points in Space," The Philosophical Magazine, 1901.
- [3]. "11 Different Uses of Dim. Reduc." <https://tinyurl.com/38k9ctnf> (02.2022).
- [4]. Giri D, Acharya UR, Martis RJ, Sree SV, Lim T-C, VI TA, and Suri JS, "Automated Diagnosis of Coronary Artery Disease Affected Patients using LDA, PCA, ICA and Discrete Wavelet Transform," Knowledge-Based Systems, 2013.
- [5]. Gumus E, Kilic N, Sertbas A, and Ucan ON, "Evaluation of Face Recognition Techniques using PCA, Wavelets and SVM," ESA, 2010.
- [6]. Jolliffe IT and Cadima J, "Principal Component Analysis: a Review and Recent Developments," Philos. Trans. R. Soc. A, 2016.
- [7]. Martis RJ, Acharya UR, and Min LC, "ECG Beat Classification using PCA, LDA, ICA and Discrete Wavelet Transform," BSSC, 2013.
- [8]. Pasini G, "Principal Component Analysis for Stock Portfolio Management," Inter. Journal of Pure and Applied Mathematics, 2017.
- [9]. Price AL, Patterson NJ, Plenge RM, Weinblatt ME, Shadick NA, and Reich D, "Principal components analysis corrects for stratification in genome-wide association studies," Nature genetics, 2006.
- [10]. Vozalis MG and Margaritis KG, "A Recommender System using Principal Component Analysis," in 11th panhellenic Conf. in informatics, 2007.
- [11]. "Dimensionality Reduction using PCA on Multivariate Timeseries Data," <https://tinyurl.com/3czux5e8>, (02.2022).
- [12]. "The PCA [...] with Time-Series," <https://tinyurl.com/54vc3cdx>, (02.2022).
- [13]. "Large-Scale [...] Detection," <https://tinyurl.com/2p8u9fv3>, (02.2022).
- [14]. Jing C and Hou J, "SVM and PCA Based Fault Classification Approaches for Complicated Industrial Process," Neurocomputing, 2015.

- [15]. Hinton GE and Salakhutdinov RR, "Reducing the Dimensionality of Data with Neural Networks," *Science*, 2006.
- [16]. Roweis ST and Saul LK, "Nonlinear Dimensionality Reduction by Locally Linear Embedding," *Science*, 2000.
- [17]. Fodor IK, "A Survey of Dimension Reduction Techniques," *Tech. re*, 2002.
- [18]. Freedman ML et al. , "Assessing the Impact of Population Stratification on Genetic Association Studies," *Nature genetics*, 2004.
- [19]. Hie B, Peters J, Nyquist SK, Shalek AK, Berger B, and Bryson BD, "Computational Methods for Single-Cell RNA Sequencing," *Annual Review of Biomedical Data Science*, 2020.
- [20]. Sav S, Bossuat J-P, Troncoso-Pastoriza JR, Claassen M, and Hubaux J-P, "Privacy-Preserving Federated Neural Network Learning for Disease-Associated Cell Classification," *Patterns*, 2022.
- [21]. Bouwmans T, Javed S, Zhang H, Lin Z, and Otazo R, "On the Applications of Robust PCA in Image and Video Processing," *IEEE*, 2018.
- [22]. Wilaiprasitporn T, Dittthaporn A, Matchaparn K, Tongbuasirilai T, Banluesombatkul N, and Chuangsuwanich E, "Affective EEG-based Person Identification using the Deep Learning Approach," *IEEE TCDS*, 2019.
- [23]. "The Ambitious Effort to Piece Together America's Fragmented Health Data," <https://tinyurl.com/mtseecfk>, (02.2022).
- [24]. Froelicher D, Troncoso-Pastoriza JR, Pyrgelis A, Sav S, Sousa JS, Bossuat J-P, and Hubaux J-P, "Scalable Privacy-Preserving Distributed Learning," *PETS*, 2021.
- [25]. Mohassel P and Zhang Y, "SecureML: A System for Scalable Privacy-Preserving Machine Learning," in *IEEE S&P*, 2017.
- [26]. Sav S, Pyrgelis A, Troncoso-Pastoriza JR, Froelicher D, Bossuat J-P, Sousa JS, and Hubaux J-P, "POSEIDON: Privacy-Preserving Federated Neural Network Learning," *NDSS*, 2021.
- [27]. Zheng W, Popa RA, Gonzalez JE, and Stoica I, "Helen: Maliciously Secure Cooperative Learning for Linear Models," in *IEEE S&P*, 2019.
- [28]. Froelicher D, Troncoso-Pastoriza JR, Raisaro JL, Cuendet MA, Sousa JS, Cho H, Berger B, Fellay J, and Hubaux J-P, "Truly privacy-preserving federated analytics for precision medicine with multiparty homomorphic encryption," *Nature communications*, vol. 12, no. 1, pp. 1–10, 2021.
- [29]. Scheibner J, Raisaro JL, Troncoso-Pastoriza JR, Ienca M, Fellay J, Vayena E, and Hubaux J-P, "Revolutionizing Medical Data Sharing Using Advanced Privacy Enhancing Technologies: Technical, Legal and Ethical Synthesis," *J Med Internet Res*, 2021.
- [30]. Keller M, Pastro V, and Rotaru D, "Overdrive: Making SPDZ Great Again," in *Eurocrypt*, 2018.
- [31]. Bogdanov D, Laur S, and Willemson J, "Sharemind: A Framework for Fast Privacy-Preserving Computations," in *ESORICS*, 2008.
- [32]. Hastings M, Hemenway B, Noble D, and Zdancewic S, "Sok: General Purpose Compilers for Secure Multi-Party Computation," in *IEEE S&P*, 2019.
- [33]. Abu-Khzam FN, Samatova NF, Ostrouchov G, Langston MA, and Geist A, "Distributed Dimension Reduction Algorithms for Widely Dispersed Data." in *IASTED PDCS*, 2002.
- [34]. Bai Z-J, Chan RH, and Luk FT, "Principal Component Analysis for Distributed Data Sets with Updating," in *APPT*, 2005.
- [35]. Balcan M-F, Kanchanapally V, Liang Y, and Woodruff D, "Improved Distributed Principal Component Analysis," *Tech. Rep*, 2014.
- [36]. Cheung Y-M and Yu F, "Federated-PCA on Vertical-Partitioned Data," *Tech. Rep*, 2020.
- [37]. Fan J, Wang D, Wang K, and Zhu Z, "Distributed Estimation of Principal Eigenspaces," *Annals of statistics*, 2019.
- [38]. Fellus J, Picard D, and Gosselin P-H, "Dimensionality Reduction in Decentralized Networks by Gossip Aggregation of Principal Components Analyzers," in *ESANN*, 2014.
- [39]. Gang A, Raja H, and Bajwa WU, "Fast and Communication-efficient Distributed PCA," in *ICASSP*, 2019.
- [40]. Liang Y, Balcan M-F, and Kanchanapally V, "Distributed PCA and K-means Clustering," in *The Big Learning Workshop at NIPS*, 2013.

- [41]. Liang Y, Balcan M-FF, Kanchanapally V, and Woodruff D, "Improved Distributed Principal Component Analysis," in NeurIPS, 2014.
- [42]. Qi H, Wang T-W, and Birdwell JD, "Global Principal Component Analysis for Dimensionality Reduction in Distributed Data Mining," Statistical data mining and knowledge discovery, 2004.
- [43]. Won H-S, Kim S-P, Lee S, Choi M-J, and Moon Y-S, "Secure Principal Component Analysis in Multiple Distributed Nodes," SCN, 2016.
- [44]. Bogdanov D, Kamm L, Laur S, and Sokk V, "Implementation and Evaluation of an Algorithm for Cryptographically Private Principal Component Analysis on Genomic Data," IEEE/ACM TCBB, 2018.
- [45]. Cho H, Wu DJ, and Berger B, "Secure genome-wide association analysis using multiparty computation," Nature biotechnology, 2018.
- [46]. Fan X, Wang G, Chen K, He X, and Xu W, "PPCA: Privacy-Preserving Principal Component Analysis Using MPC," arXiv, 2021.
- [47]. Halko N, Martinsson P-G, and Tropp JA, "Finding Structure with Randomness: Probabilistic Algorithms for Constructing Approximate Matrix Decompositions," SIAM review, 2011.
- [48]. Mouchet C, Troncoso-pastoriza JR, Bossuat J-P, and Hubaux JP, "Multiparty Homomorphic Encryption from Ring-Learning-with-Errors," in PETS, 2021.
- [49]. Kairouz P et al. , "Advances and Open Problems in Federated Learning," Foundations and Trends R in Machine Learning, 2021.
- [50]. Asharov G, Jain A, López-Alt A, Tromer E, Vaikuntanathan V, and Wichs D, "Multiparty Computation with Low Communication, Computation and Interaction via Threshold FHE," in Eurocrypt, 2012.
- [51]. Cramer R, Damgård I, and Nielsen JB, "Multiparty Computation from Threshold Homomorphic Encryption," in Eurocrypt, 2001.
- [52]. López-Alt A, Tromer E, and Vaikuntanathan V, "Cloud-Assisted Multiparty Computation from Fully Homomorphic Encryption," ePrint, 2011.
- [53]. Mouchet C, Bertrand E, and Hubaux J-P, "An Efficient Threshold Access-Structure for RLWE-Based Multiparty Homomorphic Encryption," ePrint, 2022.
- [54]. LeCun Y and Cortes C, "MNIST Handwritten Digit Database," <http://yann.lecun.com/exdb/mnist/>, 2010.
- [55]. "GWAS of Lung Cancer Susceptibility in Never-Smoking Women in Asia," <https://tinyurl.com/mr4b7n5y>, (03.2022).
- [56]. Juvekar C, Vaikuntanathan V, and Chandrakasan A, "GAZELLE: A low latency framework for secure neural network inference," in USENIX, 2018.
- [57]. Halevi S and Shoup V, "Algorithms in HELib," in CRYPTO, 2014.
- [58]. Halevi S and Shoup V, "Faster Homomorphic Linear Transformations in HELib," in CRYPTO, 2018.
- [59]. Jiang X, Kim M, Lauter K, and Song Y, "Secure Outsourced Matrix Computation and Application to Neural Networks," in ACM CCS, 2018.
- [60]. Mishra PK, Rathee D, Duong DH, and Yasuda M, "Fast Secure Matrix Multiplications over Ring-based Homomorphic Encryption," ISP, 2021.
- [61]. Cheon JH, Kim A, Kim M, and Song Y, "Homomorphic Encryption for Arithmetic of Approximate Numbers," in ASIACRYPT, 2017.
- [62]. Fan J and Vercauteren F, "Somewhat Practical Fully Homomorphic Encryption." IACR Cryptology, 2012.
- [63]. Desmedt YG, "Threshold cryptography," ETT, 1994.
- [64]. Zhu R, Ding C, and Huang Y, "Practical MPC+FHE with Applications in Secure Multi-Party Neural Network Evaluation," ePrint, 2020.
- [65]. Kim T, Kwak H, Lee D, Seo J, and Song Y, "Asymptotically Faster Multi-Key Homomorphic Encryption from Homomorphic Gadget Decomposition," ePrint, 2022.
- [66]. Kwak H, Lee D, Song Y, and Wagh S, "A Unified Framework of Homomorphic Encryption for Multiple Parties with Non-Interactive Setup," ePrint, 2021.

- [67]. Pereira HV and Aranha DF, "Principal Component Analysis over Encrypted Data using Homomorphic Encryption," in WAHC, 2016.
- [68]. De Sa C, He B, Mitliagkas I, Ré C, and Xu P, "Accelerated Stochastic Power Iteration," PMLR, 2018.
- [69]. Garber D, Shamir O, and Srebro N, "Communication-Efficient Algorithms for Distributed Stochastic Principal Component Analysis," arXiv, 2017.
- [70]. Hartebrød A, Röttger R, and Blumenthal DB, "Federated Singular Value Decomposition for High Dimensional Data," arXiv, 2022.
- [71]. Wu SX, Wai H-T, Li L, and Scaglione A, "A Review of Distributed Algorithms for Principal Component Analysis," IEEE, 2018.
- [72]. Nasr M, Shokri R, and Houmansadr A, "Comprehensive privacy analysis of deep learning: Passive and active white-box inference attacks against centralized and federated learning," in IEEE S&P, 2019.
- [73]. Melis L, Song C, De Cristofaro E, and Shmatikov V, "Exploiting Unintended Feature Leakage in Collaborative Learning," in IEEE S&P, 2019.
- [74]. Liu Y, Chen C, Zheng L, Wang L, Zhou J, Liu G, and Yang S, "Privacy Preserving PCA for Multiparty Modeling," arXiv, 2020.
- [75]. Grammenos A, Mendoza-Smith R, Mascolo C, and Crowcroft J, "Federated Principal Component Analysis," NeurIPS, 2019.
- [76]. Imtiaz H and Sarwate AD, "Differentially Private Distributed Principal Component Analysis," in ICASSP, 2018.
- [77]. Wang D and Xu J, "Principal Component Analysis in the Local Differential Privacy Model," Theoretical Computer Science, 2020.
- [78]. Charikar M, Chen K, and Farach-Colton M, "Finding Frequent Items in Data Streams," inICALP, 2002.
- [79]. Wang T-L, "Convergence of the Tridiagonal QR Algorithm," Linear algebra and its applications, 2001.
- [80]. Lyubashevsky V, Peikert C, and Regev O, "On Ideal Lattices and Learning with Errors over Rings," in EUROCRYPT, 2010.
- [81]. Costache A, Curtis BR, Hales E, Murphy S, Ogilvie T, and Player R, "On the Precision Loss in Approximate HE," ePrint, 2022.
- [82]. Lee Y, Lee J-W, Kim Y-S, Kim Y, No J-S, and Kang H, "High-Precision Bootstrapping for Approximate Homomorphic Encryption by Error Variance Minimization," in Eurocrypt, 2022.
- [83]. Kim A, Papadimitriou A, and Polyakov Y, "Approximate Homomorphic Encryption with Reduced Approximation Error," in RSA Conference, 2022.
- [84]. Li B and Micciancio D, "On The Security of Homomorphic Encryption on Approximate Numbers," in Eurocrypt, 2021.
- [85]. Cheon JH, Hong S, and Kim D, "Remark on the Security of CKKS Scheme in Practice." IACR Cryptol, 2020.
- [86]. de Castro L, Juvekar C, and Vaikuntanathan V, "Fast Vector Oblivious Linear Evaluation from Ring Learning with Errors." IACR Cryptol, 2020.
- [87]. Shanks D, "Class Number, a Theory of Factorization, and Genera," in Proc. of Symp. Math. Soc, 1971, 1971.
- [88]. Householder AS, "Unitary Triangularization of a Nonsymmetric Matrix," JACM, 1958.
- [89]. Chebyshev PL, Theorie des mecanismes connus sous le nom de parallelogrammes. Imprimerie de l'Academie imperiale des sciences, 1853.
- [90]. Han K and Ki D, "Better Bootstrapping for Approximate Homomorphic Encryption," in CT-RSA, 2020.
- [91]. Ortega JM and Kaiser HF, "The LL T and QR methods for Symmetric Tridiagonal Matrices," The Computer Journal, 1963.
- [92]. "Go Programming Language," <https://golang.org>, (07.2022).
- [93]. "Lattigo," <https://github.com/tuneinsight/lattigo>, (10.2022).

- [94]. “Cothority network library,” <https://github.com/dedis/onet>, (07.2022).
- [95]. “Mininet,” <http://mininet.org>, (07.2022).
- [96]. “The Inter. Genome Sample Resource,” <https://www.internationalgenome.org/>, (03.2022).
- [97]. Pedregosa F et al. , “Scikit-learn: Machine Learning in Python,” Journal of ML Research, 2011.
- [98]. “Wine Quality,” <https://tinyurl.com/4fk83ezw>, (03.2022).
- [99]. “Pima Indians Diabetes Dataset,” <https://tinyurl.com/y8o3x8me>, (03.2022).
- [100]. “Statlog Vehicle Silhouettes,” <https://tinyurl.com/3eet88f2>, (03.2022).
- [101]. Beaver D, “Efficient Multiparty Protocols Using Circuit Randomization,” in CRYPTO, 1991.
- [102]. “Client Software for an End-to-End MPC Protocol for PCA,” <https://github.com/hhcho/smc-pca>, (03.2022).
- [103]. Shamir A, “How to Share a Secret,” Communications of the ACM, 1979.
- [104]. Froelicher D, Troncoso-Pastoriza JR, Sousa JS, and Hubaux J, “Drynx: Decentralized, Secure, Verifiable System for Statistical Queries and Machine Learning on Distributed Datasets,” IEEE TIFS, 2020.
- [105]. Lindell Y, “How to Simulate It—A Tutorial on the Simulation Proof Technique,” in Tutorials on the Foundations of Cryptography, 2017.
- [106]. Lindner R and Peikert C, “Better Key Sizes (and Attacks) for LWE-Based Encryption,” in CT-RSA, 2011.
- [107]. Dahl M, Ning C, and Toft T, “On Secure Two-Party Integer Division,” in FC, 2012.
- [108]. Catrina O and Saxena A, “Secure Computation with Fixed-Point Numbers,” in FC, 2010.
- [109]. Markstein P, “Software Division and Square Root using Goldschmidt’s Algorithms,” in RNC’6, 2004.
- [110]. Nishide T and Ohta K, “Multiparty Computation for Interval, Equality, and Comparison Without Bit-Decomposition Protocol,” in PKC, 2007.
- [111]. “Anomaly Detection in Time Series Sensor Data,” <https://tinyurl.com/2p8sk8yz> (02.2022).

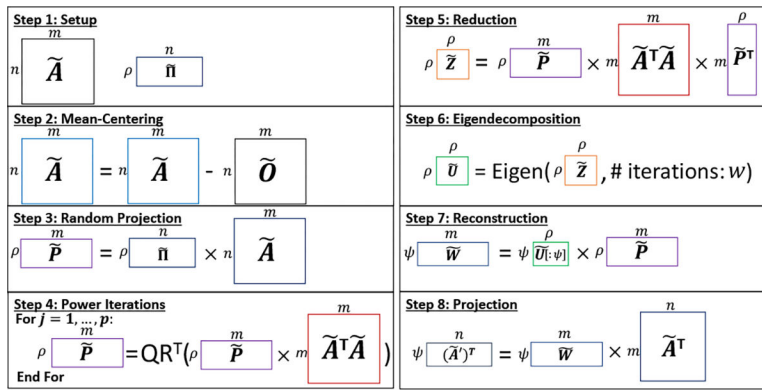


Figure 1: Randomized PCA Workflow.

Matrix dimensions are shown with the box sizes and are indicated on the left and top of the corresponding box.

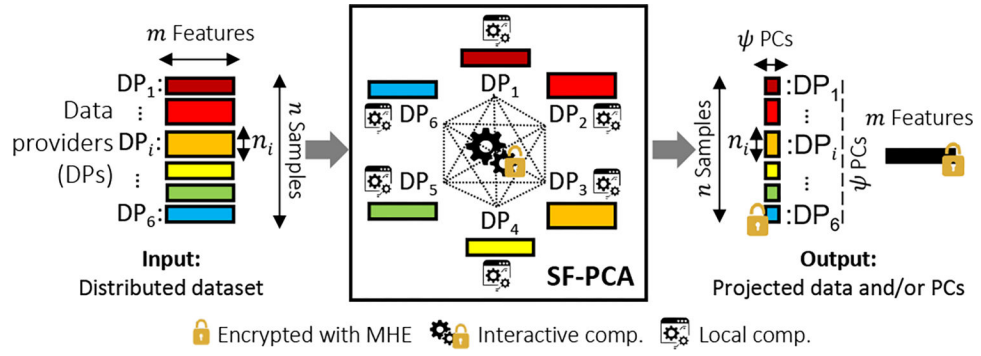


Figure 2: SF-PCA System Model and Functionality.
 Each DP_i holds a $n_i \times m$ submatrix as input and collaboratively executes SF-PCA to obtain encrypted PCs and/or the projection of its local data.

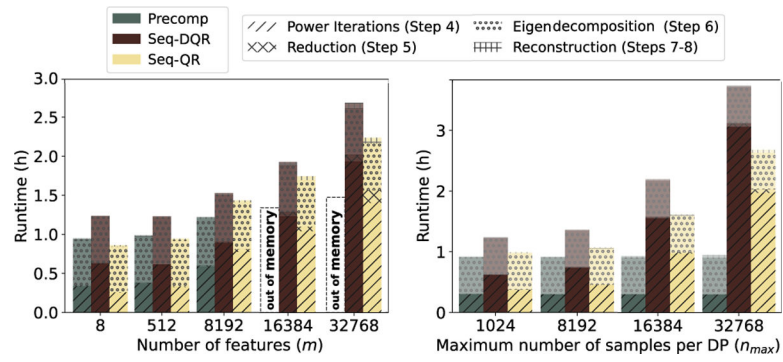


Figure 3: Runtime scaling with the number of features and samples.




Figure 4: Runtime with the number of DPs and components.

	End-to-end confidentiality	No single-trusted DP	Locality of private data	Dataset ($n \times m, p, \psi, \alpha, w$)					
				MSE; r^2 ; Runtime (hours)					
				Wine (4,898 × 11, 15, 5, 5, 1) [98]	Lung (9,098 × 23,724, 15, 2, 5, 5) [55]	Pima (767 × 8, 5, 1, 1, 5) [99]	Chr20 (2,502 × 1,773, 20, 2, 4, 1) [96]	MNIST (60,000 × 760, 40, 5, 2, 5) [54]	Vehicle (435 × 19, 20, 8, 4, 5) [100]
RPCA	✗	✗	✗	$< 10^{-6}$; 0.99; -	0.05; 0.99; -	$< 10^{-11}$; 0.99; -	$< 10^{-6}$; 0.99; -	0.3; 0.99; -	0.15; 0.99; -
Meta-an.	✗	✗	✓	$< 10^{-2}$; 0.99; -	6.55; 0.7; -	$< 10^{-2}$; 0.98; -	0.32; 0.98; -	1.48; 0.75; -	0.68; 0.99; -
SF-PCA	✓	✓	✓	$< 10^{-6}$; 0.99; 0.8	0.05; 0.99; 3.5	$< 10^{-11}$; 0.99; 0.04	$< 10^{-6}$; 0.99; 0.79	0.6; 0.91; 2.22	0.24; 0.99; 0.77
C-HE	✓	✗	✗	-; -; > 11	-; -; > 46	-; -; > 0.7	-; -; > 12	-; -; > 30	-; -; > 10
SMC	✓	✗	✗	-; -; 3.0	-; -; 9.6	-; -; 0.39	-; -; 2.32	-; -; 22.66	-; -; 6.67

Figure 5: Comparison with existing works on six real datasets.

MSE: mean-squared error, r^2 : Pearson correlation coefficient compared with ground truth PCs.

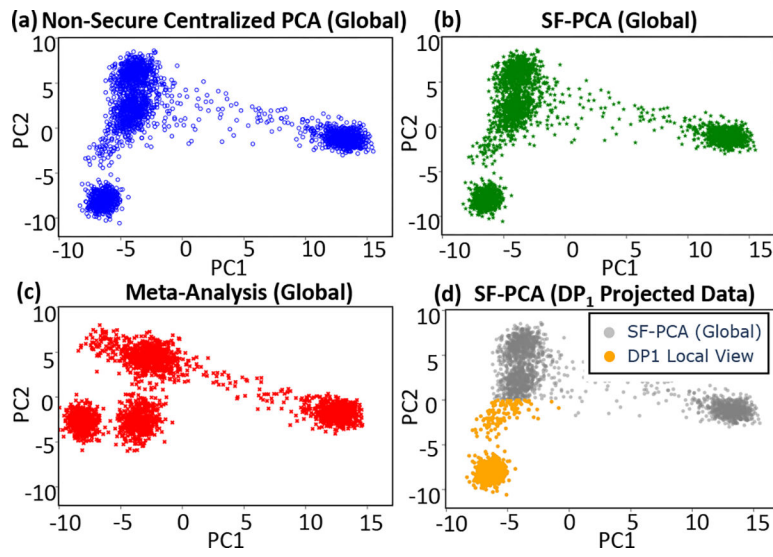


Figure 6: Demonstration of SF-PCA on Genomic Data.

TABLE 1:
Communication and computation costs of SF-PCA (Alg. 1).

$\zeta_x(\text{dim})$ returns the cost of the function x according to the dimensions (dim). The functions' costs are defined in §.6.1 for M, in Alg. 3 for QR and in Alg. 4 for Eigen.

Step	Comm.	Computation
1	$\log_2(t) + 2.5$	-
2	$1 \cdot \left\lceil \frac{m}{t} \right\rceil$	-
3	$\rho \cdot \left\lceil \frac{m}{t} \right\rceil$	-
4	$p \cdot \left(\rho \cdot \left\lceil \frac{m}{t} \right\rceil + \zeta_{\text{QR}}(\rho, m) \right)$	$p \cdot ((\text{Precomp or Seq}) + \zeta_{\text{QR}}(\rho, m)); \text{Precomp} = \zeta_M^*(\rho, m, m)$ $\text{Seq} = \zeta_M^*(\rho, m, n_i) + \zeta_M^*(\rho, n_i, m)$
5	$\rho \cdot \left\lceil \frac{\rho}{t} \right\rceil$	$\text{Precomp: } \zeta_M^*(\rho, m, m) + \zeta_M^*(\rho, m, \rho)$ $\text{Seq: } \zeta_M^*(\rho, m, n_i) + \zeta_M^*(\rho, n_i, \rho)$
6	$\zeta_{\text{Eigen}}(\rho, \rho)$	$\zeta_{\text{Eigen}}(\rho, \rho)$
7	$\psi \cdot \left\lceil \frac{m}{t} \right\rceil + \zeta_{\text{QR}}(\psi, m)$	$\text{Precomp: } \zeta_M^*(\psi, \rho, m) + \zeta_{\text{QR}}(\psi, m); \text{Seq: } \zeta_M^*(\psi, \rho, n_i) + \zeta_M^*(\psi, n_i, m) + \zeta_{\text{QR}}(\psi, m)$
8	-	$\zeta_N^*(\psi, m, n_i)$

TABLE 2:
SF-PCA's micro-benchmarks with default parameters.

Send(c) transmits a ciphertext c from one DP to another. M5 refers to the small encrypted matrix multiplication in §.6.2.

Operation	Runtime (s)	Operation	Runtime (s)
+	$7 \cdot 10^{-4}$	M1 - ($\mathbf{M}^{(8 \times 2^8)} \times \widetilde{\mathbf{N}}^{(2^8 \times 2^8)}$)	59
$\tilde{\mathbf{v}} \cdot \mathbf{c}$	0.013	M2 - ($\mathbf{M}^{(8 \times 2^8)} \times \widetilde{\mathbf{N}}^{(2^8 \times 2^8)}$)	51
$\mathbf{c} \cdot \mathbf{c}'$	0.083	M3 - ($\mathbf{M}^{(8 \times 2^8)} \times \widetilde{\mathbf{N}}^{(2^8 \times 2^8)}$)	3.8
Rot(\cdot)	0.08	M4 ($\mathbf{M}^{(8 \times 2^8)} \times \boldsymbol{\mu}^{(1 \times 2^8)}$)	0.7
$\mathbf{c}^{(1 \times 2^8)} \bullet \mathbf{c}'^{(1 \times 2^8)}$	0.73	M5 - ($\mathbf{M}^{(8 \times 8)} \times \widetilde{\mathbf{N}}^{(8 \times 8)}$)	0.9
$\text{QR}^T(\mathbf{M}^{(8 \times 2^8)})$	117	Send(c)	0.026
$\text{DQR}^T(\mathbf{M}^{(8 \times 2^8)})$	227	DBootstrap(\cdot)	0.49
$\text{QR}^T(\mathbf{M}^{(8 \times 8)})$	94	DKeyGen(\cdot)	9.0

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript