# A Privacy-Preserving Federated Learning Approach for Kernel methods

Anika Hannemann \*† anika.hannemann@informatik.uni-leipzig.de

Ali Burak Ünal<sup>‡§</sup> ali-burak.uenal@uni-tuebingen.de **Arjhun Swaminathan**<sup>‡§</sup> arjhun.swaminathan@uni-tuebingen.de

Erik Buchmann<sup>\*†</sup> buchmann@informatik.uni-leipzig.de Mete Akgün<sup>‡§</sup> mete.akguen@uni-tuebingen.de

## Abstract

It is challenging to implement Kernel methods, if the data sources are distributed and cannot be joined at a trusted third party for privacy reasons. It is even more challenging, if the use case rules out privacy-preserving approaches that introduce noise. An example for such a use case is machine learning on clinical data. To realize exact privacy preserving computation of kernel methods, we propose FLAKE, a Federated Learning Approach for KErnel methods on horizontally distributed data. With FLAKE, the data sources mask their data so that a centralized instance can compute a Gram matrix without compromising privacy. The Gram matrix allows to calculate many kernel matrices, which can be used to train kernelbased machine learning algorithms such as Support Vector Machines. We prove that FLAKE prevents an adversary from learning the input data or the number of input features under a semi-honest threat model. Experiments on clinical and synthetic data confirm that FLAKE is outperforming the accuracy and efficiency of comparable methods. The time needed to mask the data and to compute the Gram matrix is several orders of magnitude less than the time a Support Vector Machine needs to be trained. Thus, FLAKE can be applied to many use cases.

## **1** Introduction

Kernel methods are a prominent class of machine learning algorithms. However, in many real-world scenarios, kernel methods such as Support Vector Machines (SVM) cannot be readily applied, because the data sources are inherently distributed, but the data is private and cannot be shared freely. Consider a machine learning scenario, where a Kernel method on medical data is to be used to develop effective treatments, or to identify risk factors for certain diseases. The input data is collected from multiple hospitals, and it carries sensible medical information that must be kept private. In this scenario it is impossible to apply noise, because neither the patient nor the physician can accept stochastic results. The delay due to processing strong cryptography on a large data set in multiple rounds of a Secure-Multiparty Computation Protocol is also unacceptable.

<sup>\*</sup>Dept. of Computer Science, Leipzig University

<sup>&</sup>lt;sup>†</sup>Center for Scalable Data Analytics and Artificial Intelligence (ScaDS.AI) Dresden/Leipzig

<sup>&</sup>lt;sup>‡</sup>Medical Data Privacy Preserving Machine Learning (MDPPML), University of Tuebingen

<sup>&</sup>lt;sup>§</sup>Institute for Bioinformatics and Medical Informatics (IBMI), University of Tuebingen

Existing work in the field of Secure-Multiparty Computation (Mugunthan et al., 2019; Zhang et al., 2022) or Privacy-Aware Federated Learning (Pfitzner et al., 2021; Adnan et al., 2022; Malekzadeh et al., 2021) can be categorized in three approaches based on (1) encryption (2) differential privacy or (3) randomized masking (Monreale and Wang, 2016). The first two either apply strong cryptography or add noise to private data which is a severe restriction for many use cases. In this paper, we focus on the third approach using randomized masking. We present FLAKE, our Federated Learning Approach for KErnel methods. FLAKE computes the Gram matrix over distributed data sources that store horizontally partitioned data. The Gram matrix allows various kernel matrices to be computed data. Examples for such algorithms include Support Vector Machines, Gaussian processes, kernel k-means, and more. To ensure privacy, FLAKE masks the input data at the sources. FLAKE ensures that the resulting Gram matrix is exact. In order to update the Gram matrix, only a fraction of the values need to be re-calculated. Thus, inference and update are inexpensive operations. We make three contributions:

- 1. We introduce the FLAKE protocol, which allows a function party to privately compute a Gram matrix on masked input data from multiple input parties.
- 2. We prove that both the input data and the number of features is kept private, unless function party and input parties collude and share unmasked data.
- 3. We evaluate FLAKE by experiments with medical and synthetic data.

Our formal analysis and our experiments confirm that FLAKE has the potential to open up new fields of application for Kernel-based methods on horizontally partitioned data, that must be kept private, but must be analyzed with an exact approach.

Paper structure: Section 2 reviews related work, followed by a description of FLAKE in Section 3. Section 4 analyzes the privacy properties of the protocol. Section 5 contains the experimental evaluation. Finally, Section 6 concludes.

## 2 Related Work

#### 2.1 Kernel-based Methods

Kernel-based machine learning algorithms have a well-established mathematical background. They are among the well-performing machine learning algorithms and are widely utilized in various applications (Morota and Gianola, 2014; Haywood et al., 2021). They can learn non-linear patterns in the data efficiently thanks to the kernel trick: the data is represented by a set of pairwise similarity comparisons, the kernel values, instead of explicitly mapping them into higher dimensions, where linear classification can be done. To compute these kernel values, one can use several different kernel functions such as linear, polynomial, and radial basis function (RBF). Both polynomial and RBF kernels can be computed by using the kernel matrix of linear kernel, which is the Gram matrix. The Gram matrix is a positive semi-definite matrix and its entries indicate the dot product of the corresponding samples' feature vectors. Therefore, we can formulate both kernels such that they are computable by using the entries of the Gram matrix. For instance, the polynomial kernel can be written as  $k(x, y) = (x^T y + v)^p$  where  $v \leq 0$  is a trade-off parameter and  $p \in \mathbb{N}$  is the degree of the polynomial. Similarly, the RBF kernel can be formulated as  $k(x, y) = (x^T y + v)^p$ 

 $exp(-\frac{\|x^Tx - 2x^Ty + y^Ty\|^2}{2\sigma^2})$  where  $\sigma \le 0$  is the similarity adjustment parameter. In FLAKE, we will benefit from this observation to compute the desired kernel matrices from the Gram matrix.

#### 2.2 Federated Learning

Introduced by (McMahan et al., 2017), Federated Learning (FL) allows users to reap the benefits of modeling on rich yet sensitive data stored on distributed nodes. In conventional machine learning, a model  $\mathcal{M}$  is trained by the centralized data  $\mathcal{D}_{cent}$ . However, due to privacy concerns, the data is not allowed to leave the nodes. FL addresses this problem. Participating nodes  $\mathcal{N}_1, ..., \mathcal{N}_n$  in FL aim to collaboratively train the model  $\mathcal{M}$  without revealing their data to other nodes. In FL, every node  $\mathcal{N}_i$  trains a local model  $\mathcal{M}_i$  on its respective data set  $\mathcal{D}_i$  and subsequently shares the model parameters with a central server. The central server then aggregates the received model parameters to obtain a global model  $\mathcal{M}_{fed}$  with an accuracy of  $acc_{fed}$ . As more data is collected, the process is repeated,

with each node updating its local model and forwarding the updated parameters to the central server. Thus, the data does not leave its origin at any time in the computation. At some point in the iteration of FL, if  $|acc_{fed} - acc_{cent}| \leq \delta$ , then the Federated Learning framework is said to have  $\delta$ -accuracy loss. The goal in FL is to have less accuracy loss while maintaining efficiency and the data's privacy.

The privacy of the aggregated models can be ensured in different ways. Approaches based on **encryption (1)** like homomorphic encryption (HE) (Wibawa et al., 2022; Stripelis et al., 2021) aim to protect the privacy of aggregated models by encrypting individual models, but HE is computationally heavy and limited in functionality. Another cryptographic approach is secure multi-party computation (SMC) (Mugunthan et al., 2019; Zhang et al., 2022), which allows multiple parties to jointly compute on private data without revealing it, but SMC still requires significant execution time due to communication overhead.

FL studies utilizing methods based on **differential privacy (2)** (DP), protect the privacy of the aggregated model by adding noise to the individual models, making it impossible to restore the original model or to inference information about a data point's membership. However, this usually involves a cutback in accuracy (Pfitzner et al., 2021; Adnan et al., 2022; Malekzadeh et al., 2021).

The randomized masking approach (3) for FL was used by (Chen and Liu, 2005; Chen et al., 2007) who propose a geometric perturbation approach to preserve data privacy in classification tasks by hiding content while maintaining dot product and Euclidean distance relationships. To provide even stronger security, (Lin et al., 2015) utilize a random linear transformation scheme that requires the data owner to send perturbed data to the service provider for training SVM classifiers. Lin also applies perturbation for clustering tasks using a randomized kernel matrix to hide dot product and distance information (Lin, 2013). Another randomization technique using Bloom filters enables outsourcing of mining association rules while protecting business intelligence and customer privacy, but only supports approximate reconstruction of mined frequent item sets by the data owner (Qiu et al., 2008). (Yu et al., 2006) introduce random kernels where the original data gets transformed using random linear transformation. However, due to the nature of approximation and introduction of noise, they all suffer from performance loss to provide privacy. (Unal et al., 2021) provide an exact protocol and is, therefore, the closest study to our approach. Here, the data sources first have to communicate with each other to mask their data. Then they send these masked samples to enable the cloud so that it can compute the desired kernel-based machine learning algorithm. However, due to the utilized encoding technique in ESCAPED, one has to run the protocol from scratch whenever there is new data in any party that needs to be integrated into the model or there is a new party involved in the computation.

# **3** FLAKE

This section explains FLAKE, our privacy-aware Federated Learning Approach for KErnel methods.

## 3.1 Scenario Definition

We assume a multi-party scenario consisting of multiple input parties (Alice, Bob, Charlie for simplicity) and one function party. Alice, Bob and Charlie hold sensitive data that is horizontally partitioned, i.e., each input party stores the same schema with different training data. The function party performs Federated Learning iteratively on a (possibly large) set of input-data chunks. We consider a fully untrusted setting where the input data must not leave their origin. Formally, we assume an arbitrary subset of semi-honest input parties and function party where no party colludes with another one. Note that this leaves aside extreme data distributions or all-zero cases where properties of the training data of one or more input parties can be guessed, or where only one input party exist. Therefore, FLAKE needs to deal with four requirements:

*Privacy:* The function party or an input party cannot learn the data of another input party, and the number of features is kept private from the function party.

*Accuracy:* The accuracy of the federated model must be as good as that of the centralized one. *Updatability:* It must be possible to update the model with new data.

*Efficiency:* Communication costs and execution time must be feasible for our scenario.



Figure 1: Masking and Training with FLAKE: (1) Initially, each party has its own data. (2) A random matrix N and its left inverse L are computed, based on a shared seed. (3) According to protocol, the data gets masked. (4) Masked data is transferred. (5) The function party computes the dot products. (6) The Gram matrix is formed by the dot products and their transposed values. (7) A kernel matrix is computed from the Gram matrix. (8) Finally, a classifier is trained.

## 3.2 The FLAKE Protocol

FLAKE computes the Gram matrix of samples from different input parties to enable the training and testing of kernel-based machine learning algorithms. This takes place in three stages *Distribution of Seed*, *Masking and Training* and *Inference and Updating*.

**Distribution of Seed** FLAKE relies on a Public Key Infrastructure, which delivers each input party the public signing keys for all other input parties. To initiate the process, one input party is randomly selected as the leader and generates a random seed. The leader then shares this seed with the other input parties using public-key encryption and digital signatures. The function party is a natural choice for the task of the aggregator, which transmits encrypted messages between input parties, but cannot decrypt or modify these messages. We assume a trusted third party for the distribution of public keys. This is a common assumption in frameworks for privacy-preserving federated learning (Bonawitz et al., 2017; Zhang et al., 2020).

**Masking and Training** The objective of this stage is to let the function party compute a Gram matrix without learning the data from the input parties (Requirement *Privacy*). The Gram matrix G of the data matrices A, B, C provided by Alice, Bob, and Charlie is the matrix of all possible inner products  $AB^T, AC^T, CA^T, \ldots$  For better understanding, we introduce the private calculation of  $AB^T$  given  $A \in \mathbb{R}^{n_A \times f}$  and  $B \in \mathbb{R}^{n_B \times f}$  where f > 1. The following protocol reveals  $AB^T$  to the function party while hiding the input data and the number of features:

First, Alice and Bob calculate a random full-rank matrix  $N \in \mathbb{R}^{k \times f}$  for some k > f, based on a shared seed. Throughout all input parties and training iterations, N remains constant. Since rank(N) = f, there exists a non-unique matrix  $L \in \mathbb{R}^{f \times k}$  such that  $LN = I_{f \times f}$ , that can be computed using the singular value decomposition (SVD) of N, and is called the Moore-Penrose pseudoinverse. SVD allows us to write N as  $N = USV^T$  with U, V being orthogonal matrices and Sbeing a diagonal matrix. The inverse of N can be determined from the SVD:  $L = N^{-1} = US^{-1}V^T$ . Here,  $S^{-1}$  can be derived by taking the multiplicative inverse of every entry of S. Now, Alice computes a independent left inverse  $L_A$  such that  $L_A N = I$  and Bob  $L_B$  such that  $L_B N = I$ . Then, the data gets masked by Alice  $A' = AL_A(NN^T)^{\frac{1}{2}} \in \mathbb{R}^{n_A \times k}$  while Bob masks his data accordingly  $B' = BL_B(NN^T)^{\frac{1}{2}} \in \mathbb{R}^{n_B \times k}$ . Figure 1 illustrates this.

A' and B' are forwarded to the function party, which only reveals  $n_a$  and  $n_b$ , respectively, and the Gram matrix of A and B when  $A'B'^T$  is computed. The function party computes the dot product  $AB^T$ , as shown in Figure 1. Then, the function party can compute the desired kernel matrix using the Gram matrix and perform training and testing of the designated kernel-based machine learning

algorithm. The remaining entries of the Gram matrix are masked analogously. When dealing with more than three parties, the Gram matrix has to be extended correspondingly.

**Inference and Updating** To integrate new data without having to rebuild the model from scratch (Requirement *Updatability*), FLAKE provides a protocol for inference and updating the Gram matrix. We can distinguish two cases: First, one of the input parties may have received new input data. Second, a new input party shall be integrated into the computation. For simplicity, we again explain our protocol with three parties Alice, Bob and Charlie with their respective data sets A, B, C.

Assume C has new data X which must be integrated into the Gram matrix shown in Table 1. X is the data set to be used for updating the model. To extend the gram matrix with the new values from C, the function party only needs to have the entries in the dashed rectangles. The party C uses the aforementioned masking and sending approaches for this purpose. Now assume that a new input party needs to be added. In this case, the function party must calculate the values in the continuous rectangles in Table 1. The remaining new entries can be computed locally by C. In both cases, updating the Gram matrix means that the function party has to calculate only a small set of new values. The vast majority of values need to be calculated just once, and a large share of the calculation effort remains at the input parties. Note that X can be also a test data set.

When a party wants to leave the consortium the function party deletes all random components coming from this party and gram matrix entries that are calculated using these random components. This is important for compliance with legal regulations such as General Data Protection Regulation (GDPR) (European Parliament and Council of the European Union, 2016). It can be seen as an application of machine unlearning. In current FL methods, it is unclear and difficult how to eliminate a party's contribution to the collaboratively trained ML model.

Input Parties	А	В	С	Х
А	$AA^T$	$AB^T$	$AC^T$	$A\bar{X}^{\bar{T}}$
В	$BA^T$	$BB^T$	$BC^T$	$BX^T$
С	$CA^T$	$CB^T$	$CC^T$	$CX^T$
Х	$(XA^T)$	$XB^T$	$XC^T$ )	$XX^T$

Table 1: Gram matrix of three-input parties.

## 4 Analysis of Privacy Properties

## 4.1 Privacy Definition

We consider the semi-honest (or honest-but-curious) adversary model. In a multi-party scenario, a **semi-honest adversary** (Evans et al., 2018) corrupts an arbitrary subset of the parties involved. The corrupted parties follow the multi-party protocol as specified, i.e., the output of the protocol is correct. The corrupted parties try to learn private data from the messages they receive from uncorrupted parties. At the end of the protocol, the corrupted parties are allowed to share their information.

FLAKE consists of a function party and a number of input parties. From Requirement *Privacy* follows that FLAKE needs to ensure two privacy properties: (i) the data of uncorrupted input parties must kept private from any corrupted input party or the function party, and (ii) a corrupted function party must not be able to learn the number of features.

If the function party and all input parties operate honestly, privacy properties (i) and (ii) are ensured. If all input parties have been corrupted by a semi-honest adversary, privacy cannot ensured. Between these extreme cases, we distinguish three cases for further analyses:

- (1) A subset of the input parties is corrupted by a semi-honest adversary.
- (2) The function party is corrupted by a semi-honest adversary.
- (3) The function and a subset of input parties are corrupted by a semi-honest adversary.

Recall that we do not consider extreme scenarios. In particular, we exclude data distributions where the number of features or the training data of one or more input parties can be guessed, and protocols with only one input party. However, to make the guessing harder, the input parties generate a unique matrix L in each iteration. Therefore, the function party can not determine if an input party updates their data in a subsequent iteration. Also, all-zero rows are not allowed; though these are usually discarded as part of preprocessing anyway.

#### 4.2 Privacy Analysis

Before we begin analysing the privacy of the protocol, we shall establish its correctness, which is unaffected by the existence of a semi-honest adversary.

*Proof.* Without loss of generality, we assume there are two input parties Alice and Bob with individual left inverses  $L_A$  and  $L_B$  of a common mask matrix N, whose outputs are  $A' = AL_A(NN^T)^{\frac{1}{2}}$  and  $B' = BL_B(NN^T)^{\frac{1}{2}}$ . Then, the correctness of the protocol follows as below.

$$A'B'^{T} = AL_{A}(NN^{T})^{\frac{1}{2}}(BL_{B}(NN^{T})^{\frac{1}{2}})^{T},$$
  
=  $AL_{A}(NN^{T})^{\frac{1}{2}}(NN^{T})^{\frac{1}{2}}L_{B}^{T}B^{T},$   
=  $AL_{A}(NN^{T})L_{B}^{T}B^{T},$   
=  $A(L_{A}N)(L_{B}N)^{T}B^{T},$   
=  $AB^{T} = (BA^{T})^{T}.$ 

Analogously, correctness follows for  $AA^T$  and  $BB^T$ .

We analyze **Case** (1) first. Since the input parties know the number of features, we only have to prove Property (ii), i.e., a corrupted function party cannot learn the number of features.

**Theorem 4.1.** FLAKE is secure against a semi-honest adversary who corrupts a subset of the input parties.

*Proof.* Let  $S_U$  be the set of all input parties involved in the computation. While executing FLAKE protocol, an input party  $P \in S_U$  has access only to the common mask N, the common seed used to generate N and the left inverse  $L_P$  of N generated by P. At any point in FLAKE protocol, the input party P gets neither the masked data of other input parties nor the computed Gram matrix using the masked data of all input parties. Thus, A semi-honest adversary corrupting a subset of input parties  $S_C \subset S_U$  cannot learn the data of non-corrupted input parties  $S_H \subset S_U$  where  $S_C \cap S_H = \emptyset$ .

FLAKE is, therefore, secure against the semi-honest adversary corrupting a subset of input parties. Because a semi-honest adversary follows the protocol, the data provided by the corrupted input parties do not affect the result of the computation.  $\Box$ 

Regarding **Case (2)**, we need to prove that FLAKE does not allow a semi-honest function party to learn (i) input data nor (ii) the number of features.

**Theorem 4.2.** FLAKE is secure against a semi-honest adversary who corrupts the function party.

*Proof.* A semi-honest function party is only the receiver of the masked data from the input parties, and follows the protocol as intended. Without loss of generality, let there be two input parties Alice and Bob with input data  $A \in \mathbb{R}^{n_A \times f}$  and  $B \in \mathbb{R}^{n_B \times f}$ , respectively, where  $n_x$  is the number of samples in the corresponding party and f is the number of features. The semi-honest function party receives the masked input matrices of them, which are  $A' = AL_A(NN^T)^{\frac{1}{2}} \in \mathbb{R}^{n_A \times k}$  and  $B' = BL_B(NN^T)^{\frac{1}{2}} \in \mathbb{R}^{n_B \times k}$  where k > f. Then, it computes  $A'B'^T = AB^T \in \mathbb{R}^{n_A \times n_B}$ ,  $A'A'^T = AA^T \in \mathbb{R}^{n_A \times n_A}$  and  $B'B'^T = BB^T \in \mathbb{R}^{n_B \times n_B}$ . The data that the function party has access to then includes

- (a) A' and analogously, B'.
- (b)  $AB^T = (BA^T)^T$ ,  $AA^T$  and analogously  $BB^T$ .

Regarding (a), it is trivial that A' does not reveal the number of features of A. We now show that A' is not produced by a unique matrix A. Given an orthogonal matrix  $O \in \mathbb{R}^{f \times f}$  with f > 1, for  $\tilde{A} = AO$  and  $L_{\tilde{A}} = O^T L_A$ , we have  $A' = \tilde{A}(L_{\tilde{A}}(NN^T)^{\frac{1}{2}})$ . Further, since we require that not all entries of any one sample is full of zeroes, the function party cannot deduce anything about A from A'.

Regarding (b), the matrices that produce these Gram matrices are not unique, since for any orthogonal matrix  $O \in \mathbb{R}^{f \times f}$  where f > 1, labeling  $\tilde{A} = AO$  and  $\tilde{B} = BO$ , we have

$$\tilde{A}\tilde{A}^T = AA^T, \quad \tilde{B}\tilde{B}^T = BB^T, \quad \tilde{A}\tilde{B}^T = AB^T.$$

In consequence, the function party only learns the singular values and singular vectors of the matrices, i.e., it can find U and S from the singular value decomposition  $A = USV^T$  by eigen-decomposing  $AA^T$ . However, these values are insufficient to solve for A since we can generate countless number of different orthogonal matrices (Aguilera and Pérez-Aguila, 2004). The function party learns neither (i) input data nor (ii) the number of features.

Although the function party obtains the Gram matrix, it cannot deduce the samples used to compute this Gram matrix, which was shown by (Ünal et al., 2021). Details can be found in the supplementary material.

**Case (3)** means that not only the function party, but also a subset of the input parties has been corrupted by a semi-honest adversary. In this case, since the adversary knows N, the privacy of the data of the other parties is compromised since for data from a non-corrupt party Charlie of the form  $C' = CL_C(NN^T)^{\frac{1}{2}}$ , the adversary can obtain C by multiplying the data with  $(NN^T)^{\frac{1}{2}}L^T$ .

#### **5** Experiments

## 5.1 Implementation

In this section, we evaluate the performance of FLAKE and provide a run-time analysis.

We experiment with three clinical data sets which contain medical records and, thus, have strong privacy concerns (Wolberg et al., 1992; Ünal et al., 2019; Center for Machine Learning and Intelligent Systems, 2023). All of them are suitable for classification tasks. For the run-time analysis, we experimented with a synthetic data set with {500, 1000, 2000, 4000, 8000} data points (dp) for each input party. Details about their statistics can be found in the supplementary material.

Before starting with the run-time experiments, we want to compare FLAKE to other methods for randomization-based kernel computation for horizontally shared data. For this purpose, we implemented a 5-fold cross validation with FLAKE, ESCAPED (Ünal et al., 2021), PPSVM (Yu et al., 2006), RSVM (Lin et al., 2015) and a naive SVM classifier in Python. Our experiments show that FLAKE, ESCAPED and the naive classifier produces the same results as they are exact solutions. Because of the introduced stochasticity, RSVM and PPSVM have a performance almost as good as the naive classifier, but they are not exact. Furthermore, the overhead associated with the various methods was measured for a single node and 1000 data points. The overhead for all methods was found to be extremely low, to the point of being negligible. Therefore, the subsequent experiments will primarily focus on scaling up the number of data points and input parties for FLAKE and ESCAPED, the two exact methods. For further details see the supplementary material.

We implemented FLAKE for a scenario with three input parties and one function party. To mimick the network communication between input parties and function party, we have implemented each party as an isolated process that communicates with others via TCP connections. Our four data sets are divided into three disjoint partitions. Each partition is assigned an input party. Each input party then masks its data according to the FLAKE protocol, and splits the masked data into chunks. After that, each input party compresses the chunks by zlib's Deflate-algorithm, and forwards the compressed chunks to the function party. The function party deflates the chunks, computes the Gram matrix and a polynomial kernel. Finally, a SVM is trained with a 5-fold cross-validation. A grid search optimizes the corresponding hyperparameters  $C \in \{2^{-4}, ..., 2^{10}\}$  (misclassification penalty) and  $p \in \{1, ..., 5\}$  (degree).



Figure 2: Run times of FLAKE masking data sets of 1000 - 8000 dp for three input parties each. a) Time to mask the data for one input party. b) Time for computing the Gram matrix from the masked data. c) Training time of SVM

All experiments were executed on a host with an AMD 7713 with 2.0GHZ and 512 GB of memory, which is a typical stand-alone server configuration for a small datacenter. We have used a single-threaded implementation. We repeated each experiment 10 times.

#### 5.2 Run-time Analysis

We want to confirm that training time, masking time, communication time, gram-computation time and update time do not limit the applicability of FLAKE. As known from literature, SVMs typically do not scale readily to very large data sets. In a centralized scenario, it is the training time for the SVM that limits the size of the input data. We declare success, if we can show that the run-times of the stages of FLAKE in a federated scenario are negligible, compared to the stages required for the federated training of a SVM without masking.

**Training Time** The training takes place at the function party. Figure 2c shows the training time for varying numbers of dp in our synthetic data set. As expected, the longest takes the training of the data set with 8000 dp with 516.62 ( $\pm$  2.45) on average. Recall that 8000 dp means that each of our three input parties sends a masked data set of this size to the function party.

**Masking Time** To find out how much masking burdens the input parties, we ran a series of experiments, again with the synthetic data set. We varied the number of dp and measured the time for masking. Figure 2a reports the masking time measured for one input party. Even with 8000 dp per input party, the execution takes less than  $0,003 (\pm 0.0001)$  seconds on average. This masking time is negligible, compared to the time to train the SVM model, and does not restrict the applicability of FLAKE.

**Communication Time** Because our implementation runs on a data-center host, we estimate the communication time needed to send masked data from the input parties to the function party. The communication time T can be estimated as shown in Equation 1:

$$T = \frac{\text{Datasize}}{\text{Bandwidth}} + \text{Latency} \times (1 + \text{Packetloss})$$
(1)

Our largest data set consists of 8000 points, which adds up to a Datasize of 1.31 MB for each input party. A typical VPN has a Bandwidth of 1.25MBps, with an average Latency of 0.1s and a Packetloss of 2% (Ookla, 2022). For this set of parameters, the estimated the communication time is 1.05 seconds. Without Latency and Packetloss, it is 1.048 seconds. Recall that our experiments are executed on a single data-center host, i.e., the actual data transfer takes place as inter-process communication in the main memory of the host and virtually requires no time.

**Gram-Computation Time** We also measured the time the function party needs to compute the Gram matrix from the masked data from the input parties. Figure 2b shows that the computation time increases slightly more than linearly with the size of the data set, with no outliers. For 8000 dp, it took 0.99 ( $\pm$  0.0083) seconds on average to compute the Gram matrix. Again, 8000 dp means the function party receives 3x8000 masked data sets from our three input parties. In summary, we have confirmed that the Gram-computation time does not contribute much to the total computation time.

**Update Time** Having shown that the time required to mask the data, send them to the function party, and compute the Gram matrix is several orders of magnitude below the time to train the model, we now consider updating the model. To mimick a typical Federated Learning use case, where the training data increases due to dynamic data collection after the initial training, the data sets were updated with additional data in multiple training iterations.

In particular, we performed multiple training iterations starting with a synthetic data set with 1000 dp for each input party. Figure 3 reports the run-times for masking the data and computing the Gram matrix for a three party scenario. We compared four training iterations of FLAKE and ESCAPED (Ünal et al., 2021), where 1000 dp are added in each iteration. The experiment is measured in the same way as for the other diagrams. The figure confirms that FLAKE outperforms ESCAPED. In particular, masking with ESCAPED takes much more time. We conclude that updating the training data in FLAKE is an inexpensive operation and, thus, can be successfully applied in a FL setting.



Figure 3: Run-times (s) for calculation of Gram matrix (red) and for masking of data (blue) with FLAKE (F) and ESCAPED (E).

#### 5.3 Discussion

Many privacy-preserving machine learning methods ensure privacy by adding stochasticity, which decreases the result quality (privacy  $\sim$  utility trade-off) (Chen and Liu, 2005; Chen et al., 2007; Lin et al., 2015; Lin, 2013). In contrast, the function party in FLAKE obtains an exact Gram matrix (Requirement *Accuracy*), that can be used to compute any desired kernel matrix and later train any kernel-based machine learning algorithm, as if it was centralized data. ESCAPED, which provides an accurate solution as well, requires more communication between the parties, which results in longer execution times (Ünal et al., 2021). As shown in section 5, FLAKE is more efficient due to less communication rounds. Also, FLAKE allows input parties to update the Gram matrix with new samples independently of the previous samples. In ESCAPED, updating the Gram matrix with new samples is not supported. Instead, the Gram matrix must be recomputed using all the samples that the input parties have. After all, FLAKE has various advantages over preceding work using the randomized masking approach.

#### 6 Conclusion

Federated learning is an essential aspect of distributed machine learning, particularly when data privacy is a primary concern. However, when implementing both Federated Learning and privacypreserving methods, the quality of model training can suffer as a result. In this work, we have proposed FLAKE, a Federated Learning Approach for KErnel methods, as a solution to that challenge. Our approach allows for the efficient and private computation of the Gram matrix from data that is distributed on multiple sources, enabling the training of kernel-based machine learning algorithms without any trade-offs in utility. Initially, four requirements were formulated, of which we have shown that FLAKE satisfies them: *Privacy*, *Accuracy*, *Updatability* and *Efficiency*. We showed, that FLAKE is both correct and private with regard to the considered threat models. We conducted various experiments on benchmark data sets to show FLAKE meets the accuracy and correctness of centralized models. Besides conducting experiments on well-known data sets, we also replicated the experiments of (Ünal et al., 2019) on HIV V3 Loop Sequence data. While other privacy-preserving techniques can be computationally expensive, FLAKE is quite efficient. An analysis of FLAKE and comparable approaches shows, that FLAKE is not as computationally expensive. In order to expand the capabilities of the framework, additional common machine learning operations could be incorporated as future developments. Also, the masking and processing of vertically shared data could be included in FLAKE.

We believe that FLAKE has the potential to improve healthcare outcomes and reduce costs while addressing the privacy concerns associated with machine learning on clinical data. We also think that it may find many use cases in other application domains that handle sensitive, distributed data.

## References

- Adnan, M., Kalra, S., Cresswell, J. C., Taylor, G. W., and Tizhoosh, H. R. (2022). Federated learning and differential privacy for medical image analysis. *Scientific reports*, 12(1):1–10.
- Aguilera, A. and Pérez-Aguila, R. (2004). General n-dimensional rotations.
- Bonawitz, K., Ivanov, V., Kreuter, B., Marcedone, A., McMahan, H. B., Patel, S., Ramage, D., Segal, A., and Seth, K. (2017). Practical secure aggregation for privacy-preserving machine learning. In *proceedings of the* 2017 ACM SIGSAC Conference on Computer and Communications Security, pages 1175–1191.
- Center for Machine Learning and Intelligent Systems (2023). UCI Machine Learning Repository. http://archive.ics.uci.edu.
- Chen, K. and Liu, L. (2005). Privacy preserving data classification with rotation perturbation. In *Fifth IEEE International Conference on Data Mining (ICDM'05)*, pages 4–pp. IEEE.
- Chen, K., Sun, G., and Liu, L. (2007). Towards attack-resilient geometric data perturbation. In proceedings of the 2007 SIAM international conference on Data mining, pages 78–89. SIAM.
- European Parliament and Council of the European Union (2016). Regulation (eu) 2016/679 of the european parliament and of the council of 27 april 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing directive 95/46/ec (general data protection regulation).
- Evans, D., Kolesnikov, V., Rosulek, M., et al. (2018). A pragmatic introduction to secure multi-party computation. Foundations and Trends in Privacy and Security, 2(2-3):70–246.
- Haywood, A. L., Redshaw, J., Hanson-Heine, M. W., Taylor, A., Brown, A., Mason, A. M., Gaertner, T., and Hirst, J. D. (2021). Kernel methods for predicting yields of chemical reactions. *Journal of Chemical Information and Modeling*.
- Lin, K.-P. (2013). Privacy-preserving kernel k-means outsourcing with randomized kernels. In 2013 IEEE 13th International Conference on Data Mining Workshops, pages 860–866. IEEE.
- Lin, K.-P., Chang, Y.-W., and Chen, M.-S. (2015). Secure support vector machines outsourcing with random linear transformation. *Knowledge and Information Systems*, 44:147–176.
- Malekzadeh, M., Hasircioglu, B., Mital, N., Katarya, K., Ozfatura, M. E., and Gündüz, D. (2021). Dopamine: Differentially private federated learning on medical data. *arXiv preprint arXiv:2101.11693*.
- McMahan, B., Moore, E., Ramage, D., Hampson, S., and y Arcas, B. A. (2017). Communication-efficient learning of deep networks from decentralized data. In *Artificial intelligence and statistics*, pages 1273–1282. PMLR.
- Monreale, A. and Wang, W. H. (2016). Privacy-preserving outsourcing of data mining. In 2016 IEEE 40th Annual Computer Software and Applications Conference (COMPSAC), volume 2, pages 583–588. IEEE.
- Morota, G. and Gianola, D. (2014). Kernel-based whole-genome prediction of complex traits: a review. *Frontiers in genetics*, 5:363.
- Mugunthan, V., Polychroniadou, A., Byrd, D., and Balch, T. H. (2019). Smpai: Secure multi-party computation for federated learning. In *Proceedings of the NeurIPS 2019 Workshop on Robust AI in Financial Services*.
- Ookla (2022). Speedtest global index. https://www.speedtest.net/global-index. Accessed: 01 20, 2022.
- Pfitzner, B., Steckhan, N., and Arnrich, B. (2021). Federated learning in a medical context: a systematic literature review. ACM Transactions on Internet Technology (TOIT), 21(2):1–31.
- Qiu, L., Li, Y., and Wu, X. (2008). Protecting business intelligence and customer privacy while outsourcing data mining tasks. *Knowledge and Information Systems*, 17(1):99–120.
- Stripelis, D., Saleem, H., Ghai, T., Dhinagar, N., Gupta, U., Anastasiou, C., Ver Steeg, G., Ravi, S., Naveed, M., Thompson, P. M., et al. (2021). Secure neuroimaging analysis using federated learning with homomorphic encryption. In *17th International Symposium on Medical Information Processing and Analysis*, volume 12088, pages 351–359. SPIE.

- Ünal, A. B., Akgün, M., and Pfeifer, N. (2019). A framework with randomized encoding for a fast privacy preserving calculation of non-linear kernels for machine learning applications in precision medicine. In *International Conference on Cryptology and Network Security*, pages 493–511. Springer.
- Ünal, A. B., Akgün, M., and Pfeifer, N. (2021). Escaped: Efficient secure and private dot product framework for kernel-based machine learning algorithms with applications in healthcare. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 9988–9996.
- Wibawa, F., Catak, F. O., Kuzlu, M., Sarp, S., and Cali, U. (2022). Homomorphic encryption and federated learning based privacy-preserving cnn training: Covid-19 detection use-case. In *Proceedings of the 2022 European Interdisciplinary Cybersecurity Conference*, pages 85–90.
- Wolberg, W. H., Street, W. N., and Mangasarian, O. L. (1992). Breast cancer wisconsin (diagnostic) data set. UCI Machine Learning Repository [http://archive. ics. uci. edu/ml/].
- Yu, H., Jiang, X., and Vaidya, J. (2006). Privacy-preserving svm using nonlinear kernels on horizontally partitioned data. In *Proceedings of the 2006 ACM symposium on Applied computing*, pages 603–610.
- Zeng, Z.-Q., Yu, H.-B., Xu, H.-R., Xie, Y.-Q., and Gao, J. (2008). Fast training support vector machines using parallel sequential minimal optimization. In 2008 3rd international conference on intelligent system and knowledge engineering, volume 1, pages 997–1001. IEEE.
- Zhang, C., Ekanut, S., Zhen, L., and Li, Z. (2022). Augmented multi-party computation against gradient leakage in federated learning. *IEEE Transactions on Big Data*.
- Zhang, C., Li, S., Xia, J., Wang, W., Yan, F., and Liu, Y. (2020). Batchcrypt: Efficient homomorphic encryption for cross-silo federated learning. In *Proceedings of the 2020 USENIX Annual Technical Conference (USENIX ATC 2020)*.

## 7 Supplementary Material

## 7.1 Privacy Proof

The following proof is based on a proof by Ünal et al. (2019).

**Theorem 7.1.** FLAKE provides security against a malicious function party A, assuming A is either semi-honest or malicious and does not collude with any input parties. In this scenario, A is unable to deduce the data of the input parties from the Gram matrix G that is generated as a result.

*Proof.* Although the number of features are hidden by FLAKE, we assume now the full Gram matrix  $G = DD^T$  with the data of the input parties D = [A, B, C] and the number of features are known to the function party. We show, that an attacker could not obtain any data since it there are multiple matrices that result in the Gram matrix.

Assume that there is a rotation matrix  $R \in \mathbb{R}^{N \times N}$  where  $N = 2(n_a + n_b + n_c)$  with  $n_x$  is the number of samples in the corresponding party. Then, there is a matrix E which can be computed by  $E = R^{-1}D$ . From that, we can say that D = RE. Then, due to the rotation property of  $R^{-1} = R^T$ , the the following holds:

$$K = D^{T}D$$
  
=  $(RE)^{T}(RE)$   
=  $E^{T}R^{T}RE$   
=  $E^{T}R^{-1}RE$   
=  $E^{T}E$ 

Since Aguilera and Pérez-Aguila (2004) showed, that countless rotation matrices can be generated, we cannot obtain a unique matrix resulting in Gram matrix G: For every new rotation matrix  $\theta \in \mathbb{R}^{N \times N}$ , there exists a new matrix  $\beta = \theta^{-1}D$  satisfying  $G = \beta^T \beta$ . Thus, A is unable to deduce the input parties' data D = (A, B, C) from  $G = D^D T$ .

## 7.2 Supplementary Experiments

All methods employed a polynomial kernel and identical hyperparameter settings. For this implementation, Sequential Minimal Optimization (libsvm) provided by scikit-learn was used Zeng et al. (2008). Since the Pima Indian diabetes data set, HIV and Breast Cancer data set have an unbalanced distribution of classes, we have applied Macro Averaging. Correspondingly, for the balanced synthetic data set, Micro Averaging.

	NAD	VE	ELAKE		
	INAL	vE	TEARE		
DATA SET	NUMBER OF DATA POINTS	NUMBER OF FEATURES	BINARY/MULTI - LABEL	DISTRIBUTION	
DIABETES	768	8	BINARY	IN-BALANCED	
CANCER	569	10	BINARY	IN-BALANCED	
HIV	766	924	BINARY	IN-BALANCED	
Synthetic	500-8000	20	MULTI CLASS	BALANCED	

Table 2: statistics of data sets used in the experiment section

Table 3: ROC AUC with standard deviation for Naive SVM, FLAKE, ESCAPED, RSVM, PPSVM on various data sets.

	NAIVE	FLAKE	ESCAPED	RSVM	PPSVM
DIABETES Cancer HIV Synthetic	$\begin{array}{c} 0.97 {\pm}~ 0.04 \\ 0.97 {\pm}~ 0.03 \\ 0.78 {\pm}~ 0.03 \\ 0.97 {\pm}~ 0.01 \end{array}$	$\begin{array}{c} 0.97 {\pm}~ 0.04 \\ 0.97 {\pm}~ 0.03 \\ 0.78 {\pm}~ 0.03 \\ 0.97 {\pm}~ 0.01 \end{array}$	$\begin{array}{c} 0.97 {\pm}~ 0.04 \\ 0.97 {\pm}~ 0.03 \\ 0.78 {\pm}~ 0.03 \\ 0.97 {\pm}~ 0.01 \end{array}$	$\begin{array}{c} 0.95 {\pm}~ 0.02 \\ 0.96 {\pm}~ 0.02 \\ 0.65 {\pm}~ 0.17 \\ 0.83 {\pm}~ 0.04 \end{array}$	$\begin{array}{c} 0.94 {\pm}~ 0.04 \\ 0.97 {\pm}~ 0.04 \\ 0.64 {\pm}~ 0.10 \\ 0.95 {\pm}~ 0.01 \end{array}$

Table 4: Overhead (Masking time + Gram time) for FLAKE, ESCAPED, RSVM, PPSVM for three input parties with 1000 dp each.

	FLAKE	ESCAPED	RSVM	PPSVM
Masking time for one IP Time to compute Gram	$0.00146 \\ 0.02071$	$1.23610 \\ 0.03156$	$0.00201 \\ 0.00530$	0.02257 0.01121