

EPIC: Efficient Private Image Classification, *Learning* from the Masters

Eleftheria Makri

imec-COSIC, KU Leuven

ABRR, Saxion University of Applied Sciences.

eleftheria.makri@esat.kuleuven.be

Nigel P. Smart

imec-COSIC, KU Leuven.

nigel.smart@kuleuven.be

Dragos Rotaru

imec-COSIC, KU Leuven

University of Bristol.

dragos.rotaru@esat.kuleuven.be

Frederik Vercauteren

imec-COSIC, KU Leuven.

frederik.vercauteren@esat.kuleuven.be

Abstract

Outsourcing an image classification task raises privacy concerns, both from the image provider’s perspective, who wishes to keep their images confidential, and from the classification algorithm provider’s perspective, who wishes to protect the intellectual property of their classifier.

We propose EPIC, an efficient private image classification system based on support vector machine (SVM) learning, which is secure against malicious adversaries. The novelty of EPIC is that it builds upon transfer learning techniques known from the Machine Learning (ML) literature and minimizes the load on the privacy-preserving part.

Our solution is based on Secure Multiparty Computation (MPC), it is 34 times faster than Gazelle (arXiv:1801.05507) –the state-of-the-art in private image classification– and it improves the total communication cost by 50 times, while achieving a 7% higher accuracy on CIFAR-10 dataset. When benchmarked for performance, while maintaining the same CIFAR-10 accuracy as Gazelle, EPIC is 700 times faster and the communication cost is reduced by 500 times.

1 Introduction

Visual object recognition is an important machine learning application, deployed in numerous real-life settings. Machine Learning as a Service (MLaaS) is becoming increasingly popular in the era of cloud computing, data mining, and knowledge extraction. Object recognition is such a machine learning task that can be provided as a cloud service. However, in most application scenarios, straightforward outsourcing of the object recognition task is not possible due to privacy concerns. Generally, the *image holder* who wishes to perform the image classification process, requires their input images to remain confidential (i.e., to not be revealed to the service provider). On the other hand, the *classification algorithm provider* wishes to commercially exploit their algorithm; hence, requires the algorithm parameters to remain confidential.

We consider an approach, which facilitates the outsourcing of the image classification task to an external classification algorithm provider, without requiring the establishment

of trust, contractually or otherwise, between the involved parties. We focus on the evaluation task (i.e., labeling a new unclassified image), and not the learning task. Our proposal is based on Secure Multiparty Computation (MPC), and allows for private image classification without revealing anything about the private images of the image holder, nor about the parameters of the classification algorithm. Unlike previous work [4, 5, 20], we can fully outsource the task at hand, in such a way that the classification algorithm provider does not need to be the same entity as the *cloud computing provider*. Although any of the involved parties (i.e., the classification algorithm provider, and the image holder) can play the role of one of the MPC servers, this is not a requirement for guaranteeing the security of our proposal. MPC allows distribution of trust to two or more parties. As long as the image holder (resp. the classification algorithm provider) trusts at least one of the MPC servers, their input images (resp. their classification algorithm parameters) remain secret.

EPIC, our privacy-preserving image classification solution, combines the techniques of transfer learning feature extraction, support vector machine (SVM) classification, and Secure Multiparty Computation. In this work we use recently developed techniques for generic image classification (within the ImageNet competition) such as transfer learning to extract powerful features. Then, the computation done in the MPC setting is minimized to only evaluate a linear function with secret shared inputs.

MPC allows a set of mutually distrusting parties to jointly compute a function on their inputs, without revealing anything about these inputs (other than what can be inferred from the function output itself). Currently, MPC allows one to compute relatively simple functions on private data; arbitrarily complex functions can be supported, but with large computational cost. This is why we focus on classification via SVM, as opposed to using more sophisticated techniques, such as Neural Networks (NNs), in the privacy-preserving domain. While the field of private image classification is shifting towards the deployment of NN-based approaches [26, 34, 43], we show that this is not necessary,

as we can achieve classification with better accuracy by improving the feature extraction techniques used. Although Convolutional Neural Networks (CNNs) are at present the most prevalent machine learning method for image classification [24], we confirm that SVMs can achieve high accuracy, as long as they are provided with good quality features. SVMs are also favored over NNs, because transforming a NN to a privacy-preserving one would result in an inefficient solution (e.g., 570 seconds for one image classification by CryptoNets [20]), given the non-linear nature of NNs.

A schematic representation of the application scenario treated by EPIC is given in Figure 1. Using additive secret sharing techniques both the *classification algorithm provider*, and the *image holder* share their inputs to the $n \geq 2$ MPC servers. Note that no information about the actual secret inputs can be gained by the individual shares alone. Thus, each MPC server learns nothing about the inputs of the two parties. The cluster of the MPC servers comprise the *cloud computing provider*, which together execute the MPC protocol to produce the final classification result. The MPC servers communicate via authenticated channels to accomplish what the protocol prescribes. Finally, the protocol completes its execution by having all MPC servers sending their share of the final classification result to the party prescribed by the protocol, who can then reconstruct the result by combining the received shares. This party can be the *image holder*, or an external *analyst*, assigned to examine the classification results, without getting access to the underlying private images. Note that the involved parties, namely the image holder, the classification algorithm provider, and potentially the analyst, may play the role of the MPC servers themselves, avoiding completely the outsourcing to the cloud provider(s). Specifically, we have implemented our solution using SPDZ [7], which was introduced by Damgård et al. [12, 13], and is based on additive secret sharing.

A key aspect of our work is how the data is processed before the MPC engine is used to perform the classification. The SVM classification is performed on so-called feature vectors, and not directly on the images. The way one determines these feature vectors not only affects accuracy, but it also has an impact on security. As shown in Figure 1 the image holder performs the feature extraction on the input image before it is passed to the secure gateway. Thus this feature extraction must not be specific to the algorithm classification provider; otherwise the extracted features could reveal information about exactly what is being classified. A classical form of feature extraction is to apply a Histogram of Oriented Gradients (HOG) feature extraction method, followed by Principal Component Analysis (PCA) dimensionality reduction. In this method the PCA components are selected depending on the classification task at hand, and so the PCA matrix reveals information about the objects being classified. In addition, this classical method is not as accurate as modern methods.

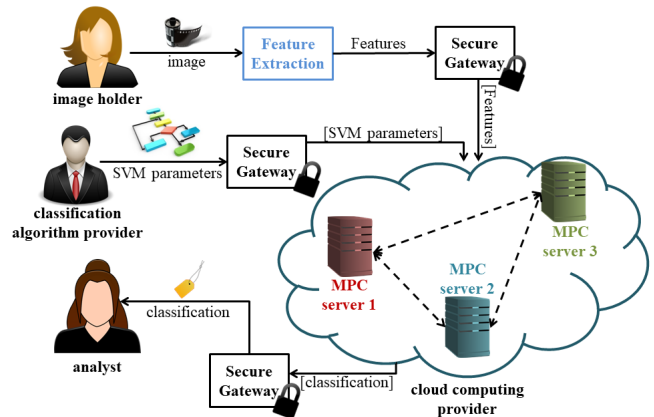


Figure 1. A schematic representation of the private image classification scenario.

Instead of this classical method we apply a generic feature extraction method, which is essentially independent of the underlying classification task. In particular we employ TensorFlow [1], to extract features based on the activation of a deep convolutional network trained on a set of object recognition tasks, *different* from the target task. The underlying CNN used for feature extraction is Inception-v3 [47]. Note that this method is known as CNN-off-the-shelf in the ML literature, and it was successfully applied in various image recognition tasks [15, 42]. The problem of transfer learning using raw CNN features was thoroughly studied by Azizpour et al. [3], and Yosinski et al. [57]. Since the CNN is generic, it can be released in the clear, and hence can be made part of the image holder’s pre-processing. This not only gives us a security benefit, but it also significantly improves the accuracy of our method.

Our proposal shows that privacy-preserving image classification has become practical. We propose two variants of EPIC, which provide different accuracy-efficiency tradeoffs. EPIC outperforms the state-of-the-art in secure neural network inference [26], both in terms of efficiency, and in terms of prediction accuracy. As shown in Table 1, we are the only work providing active security, and unlike differential privacy techniques (which add noise to the inputs to preserve privacy) our solution enjoys provable security guarantees.

Contributions: Our contributions are thus four fold: i) We enable full outsourcing of privacy-preserving image classification to a third independent party using a simple technique yet much faster and accurate than others which require complicated machinery. ii) Our solution does not leak *any* information about the private images, nor the classifier, while being the first to provide active security. iii) We show how an effective, data-independent feature extraction technique can be deployed to alleviate the privacy-preserving

computations, while increasing accuracy and computational efficiency. iv) We demonstrate the practicality of our approach, both in terms of efficiency, and in terms of accuracy, by conducting experiments on realistic datasets.

1.1 Use cases

There are a number of potential use cases for a system such as EPIC. Consider a stamp collector, who wants to sell some of their collection without revealing all about their collection, and a buyer who is interested in specific stamps. EPIC can be used by the seller to “encrypt” the images of their stamps, which are then run against a buyer provided classification model. The model simply classifies into buy-or-sell stamps, or in a more complex scenario, decides the price range of an item being sold, given its quality, as detected from the image. Thus, both parties’ inputs are protected, but matching can still be performed. This simple use case can clearly be extended to different forms of purchase, where visual inspection and selection is performed.

Another example where EPIC can be applied is targeted surveillance on cars. EPIC can be used to classify car images, where various organizations want to identify attributes such as maximum speed, car model or seat capacity [54]. In this way, EPIC maintains secret the images from a car owner perspective including the attributes the organizations are looking for.

There are many public CNNs available online in Caffe’s Model Zoo, which can be used together with our EPIC solution to add a privacy dimension to a typical ML problem [25]. In our paper we selected Inception-v3 as the public CNN to extract features, since it suits a large number of generic image recognition tasks, and allows us to benchmark EPIC against previous solutions on traditional datasets such as CIFAR-10.

1.2 Preliminaries on Secure Multiparty Computation

Secure Multiparty Computation (MPC) is a cryptographic method allowing a set of parties to jointly compute a function on their inputs, without revealing the inputs to the rest of the parties. The two main security models used to realize MPC are the *passive*, and *active* security model. In the passive security model, we assume that the protocol participants follow the protocol specification honestly, but they try to learn as much information as possible about the private inputs, during the protocol execution. In the active security model, we assume that the protocol participants may actively, and arbitrarily deviate from the protocol specification. Clearly, the active security model offers stronger security guarantees. In both models we can construct protocols that require an honest majority of the protocol participants to guarantee security, or protocols that guarantee security assuming a dishonest majority of the protocol participants. Our solution offers strong security guarantees, providing active static

security, with a dishonest majority. This means that an adversary may corrupt, prior to the protocol execution, up to $n - 1$ out of the n protocol participants, without leaking any private information, and without allowing any false protocol output to be accepted as correct.

Our solution is implemented using the SPDZ MPC framework [12, 13], and that is why it enjoys the aforementioned security properties. The computational and communication costs of the constructed protocols increase linearly in the number of protocol participants. For our experiments we assume the minimum number of MPC servers necessary to perform the outsourced computation, namely two MPC servers. Scaling to more than two servers is straightforward. We consider this to be a reasonable assumption, given that these servers can be provided by independent cloud providers, such as Google, and Amazon (and Azure if we wish to expand to three parties), who have no incentive to collude against their clients. The probability that all cloud providers get corrupted by an adversary simultaneously is small. On the other hand the two MPC servers could be the image holder and the classification algorithm holder themselves, thus removing any need for a cloud provider in the first place.

SPDZ is based on additive secret sharing, allowing the participants to share their private inputs, in such a way that no information about the private inputs is revealed to the individual participants. Additive secret sharing enjoys an additively homomorphic property, meaning that any linear function can be directly computed on the shares that each protocol participant holds, without requiring interaction amongst the parties. Upon reconstruction of the shared output (which requires all parties to send their shares of the secret), the result will be the correct result of the linear function, as if it had been applied on the secret input. Thus, we can perform additions, and multiplications with non-secret constant values on the secret shared inputs. To perform multiplications between secret shared inputs, or any other non-linear operation, we need to execute a secure interactive protocol between the MPC servers.

The SPDZ approach works in two phases: a preprocessing phase, and an online phase. The preprocessing phase can take place offline, anytime prior to the execution of the online phase. This phase only requires the MPC servers to be online, and not the inputting parties. During this phase the MPC servers create shared randomness, which the client, and the classification algorithm provider can later use to securely share their private inputs. Moreover, the MPC servers create shared random values to be consumed during the online phase, and make it efficient. For the online phase, the two inputting parties first need to provide the MPC servers with their private inputs. This is performed in a secure manner, based on the Output Delivery protocol, and Input Supply protocol, proposed by Damgård et al. [11]. Then, the MPC servers proceed with the secure computation of the actual function, as prescribed by the protocol transcript. For more

details on the MPC techniques used, we refer the reader to the work of Damgård et al. [12, 13].

2 Related Work

The related work on privacy-preserving machine learning focuses on providing a secure training phase, a secure classification phase, or both a secure training and classification. The first research works in the field aimed at designing a privacy-preserving training phase. Recently, due to the advent of cloud computing, and Machine Learning as a Service, more and more works focus on the design of a privacy-preserving classification phase. Fewer works have attempted to address both the training, and the classification phases in a privacy-preserving manner. Given the sheer number of works on private machine-learning, we only consider the works most related to ours in this section.

To facilitate an easy comparison of the related work, we summarize the main features of each proposal in Table 1.

- The first column of Table 1 is the reference to the corresponding research paper.
- The second column indicates whether the work considers a secure training phase (T), a secure training and classification phase (T+C), or only a secure classification phase (C).
- The third column indicates the security model, under which the proposed protocols are secure, where P stands for passive security, A stands for active security, and N/A (not applicable) refers to differential privacy techniques, which do not provide provable security guarantees.
- The fourth column denotes the method used to preserve privacy, where DP stands for differential privacy techniques; SP stands for selective privacy, and refers to the unique characteristic of the work of Shokri and Shmatikov [46] allowing the users to decide how much private information about their learned models they wish to reveal; MPC stands for Multiparty Computation; SHE stands for Somewhat Homomorphic Encryption; and 2-PC stands for Two-Party Computation.
- The fifth column lists the training method(s) used, where N-L SVM stands for non-linear Support Vector Machine; NN for Neural Networks; LM for Linear Means classification; FLD for Fisher’s Linear Discriminant classification; HD for hyperplane decision; LIR for linear regression, LOR for logistic regression, and DT stands for decision trees.
- The sixth column lists the information that is revealed by the protocol execution. C stands for information about the classifier; and TD stands for information about the training data. We note with boldface letters the information that is intentionally revealed by the protocol execution, and we mark with an asterisk

the information that is protected by means of non-provable security (e.g., differential privacy techniques). Information that can potentially, and unintentionally be leaked is noted with normal, non-boldface letters.

- The last column indicates whether an implementation and experimental results of the suggested method have been provided.

	Func.	Model	Privacy Mthd	Train Mthd	Info Leak	Impl.
[31]	T	N/A	DP	N-L SVM	C TD*	✓
[32]	T	N/A	DP	N-L SVM	C TD*	✓
[46]	T	P	SP	NN	C	✓
[51]	T	P	MPC	N-L SVM	C	✓
[48]	T	P	MPC DP	N-L SVM	C	✓
[9]	T	P	MPC DP	NN	C*	✓
[22]	T+C	P	SHE	LM; FLD	no	✓
[2]	T+C	P	SHE	Bayes; random forests	no	✓
[30]	T+C	P	2-PC	N-L SVM	no	×
[10]	T+C	P	2-PC	NN	TD	×
[35]	T+C	P	2-PC	NN LIR LOR	no	✓
[20]	C	P	SHE	NN	no	✓
[4]	C	P	SHE	N-L SVM	C	✓
[8]	C	P	SHE	NN	no	×
[6]	C	P	SHE 2-PC	HD; Bayes; DT	no	✓
[41]	C	P	2-PC	N-L SVM	no	✓
[5]	C	P	2-PC	NN	C	×
[36]	C	P	2-PC	NN	no	✓
[45]	C	P	2-PC	NN	no	✓
[34]	C	P	2-PC	NN	filter size	✓
[43]	C	P	2-PC	NN SVM	no	✓
[26]	C	P	2-PC	NN	no	✓
EPIC	C	A	MPC	SVM	no	✓

Table 1. Comparison of the Related Work

Training a SVM in a privacy-friendly way, has been previously considered based on techniques of differential privacy [31, 32]. Despite the little overhead that these techniques incur, which makes them competitive from an efficiency perspective, they do not provide provable security guarantees. Shokri and Shmatikov [46] achieve privacy-preserving collaborative deep learning with multiple participants, while refraining from using cryptographic techniques. Although their work focuses on learning the artificial neural network, they do consider protecting the privacy of each individual’s neural network, and allow the participants to decide how much information they wish to share about their models.

A lot of research has been devoted to *provable* privacy-preserving techniques for training a classifier. This line of research, much like ours, originates from Yao’s millionaire problem [56], describing two-party computation, and its extension to multiparty computation [21] to securely compute any generic function. Specifically, the challenge of privacy-preserving data mining has been an active research area since the seminal work of Lindell and Pinkas [33]. More recently, Vaidya et al. [51] showed how to train a SVM classifier, in a privacy-preserving manner, based on vertically, horizontally, and arbitrarily partitioned training data. In follow-up work, Teo et al. [48] improved upon the efficiency of the solution of Vaidya et al. [51], and showed that their approach scales well to address the challenges of data mining on big data. Chase et al. [9] combine MPC techniques with differential privacy techniques to achieve private neural network learning. Their work provides provable security guarantees for the learning phase (though in the passive security model), and adds noise to the final resulting network to protect its privacy.

A parallel research line aiming to address the same challenge, namely privacy-preserving data mining, is based on homomorphic encryption (instead of MPC). The notion of homomorphic encryption dates back to the work of Rivest et al. [44], but only recently fully homomorphic encryption was devised [19]. This type of homomorphic encryption allows the computation of any polynomial function on the encrypted data, and unlike MPC, it does not require communication, as the task can be outsourced to one single party. Since the seminal work of Gentry [19], a lot of somewhat homomorphic encryption schemes have been proposed, allowing computations of polynomial functions of a limited degree. Graepel et al. [22] consider both machine learning training, and classification based on encrypted data, with their solutions being secure in the passive model. Due to the selected homomorphic encryption scheme, Graepel et al. [22] cannot treat comparisons efficiently, which excludes SVM-based solutions. Addressing both learning, and classification based on extremely random forests, and naïve Bayes networks, Aslett et al. [2], also work on homomorphically encrypted data.

One of the first private SVM classifiers was proposed by Laur et al. [30], which addresses both the training and the classification in a privacy-preserving manner. Their work combines the techniques of homomorphic encryption, secret sharing, and circuit evaluation, into a passively secure 2-PC solution. Concurrently, and independently Dahl [10] is working on using the very same MPC framework as in our work, to realize both the training, and the classification phase of CNN based privacy-preserving algorithms. While Dahl [10] is deploying CNNs instead of SVM, he needs to apply them in a non-black-box fashion. The protocol of Dahl [10] allows some leakage of information during the training phase, which is not the case with our approach. SecureML [35] also considers both training and classification in the 2-PC setting, and the passive security model. These approaches [10, 30, 35] can only treat the two-party setting, and cannot be trivially extended to allow the classifier provider to be a different entity than the cloud provider.

Other works focus particularly on the private image classification problem, instead of the training of the model. Gilad-Bachrach et al. [20] propose a solution applicable to the image classification problem, based on homomorphically encrypted data. The resulting *CryptoNets* [20] provide an accuracy of 99% for the MNIST dataset, and can make on average 51739 predictions per hour. However, this is only the case when the predictions are to be made simultaneously; for a single prediction the task takes 570 seconds to complete.

Recent work by Barnett et al. [4] demonstrated the potential of polynomial-kernel SVM to be used for classification in a privacy-preserving manner. Specifically, Barnett et al. apply SVM techniques for the classification –as in our work– but on encrypted data. Although Barnett et al. mention the potential of an MPC approach to be more efficient in such a setting, they do not consider it, because direct translation of the protocols to MPC would require interaction between the client and the classification algorithm provider during the computations. We overcome this limitation by extending the application scenario in such a way that it allows the classification task to be fully outsourced to a cluster of independent third parties. We implement their approach using SPDZ in a more secure way by keeping the PCA components private (they choose to make them public). Unfortunately this method is quite expensive compared to the EPIC due to the non-linearity of the polynomial SVM and has poor accuracy as well. While inefficient and inaccurate it provides a benchmark to start with and it shows the gap between an FHE approach and an MPC based (see details in Section 4.6).

Chabanne et al. [8] attempted to approximate commonly used functions used in NN-based classification in a SHE-friendly manner. Despite the high prediction accuracy that their work achieves, Chabanne et al. do not provide any performance evaluation results.

In the 2-PC setting, Bost et al. [6], and Rahulamathavan et al. [41] focus on the problem of private classification, where

both the classifier parameters, and the client’s input to be classified need to remain private. Both approaches offer passive security, and do not consider nor experiment with linear SVM. Barni et al. [5] propose private NN-based data classification, also in the 2-PC setting and passive security model. They suggest three protocols, which offer different privacy guarantees for the classifier owner, while always protecting fully the client’s input. Follow up work by Orlandi et al. [36], also considers NN-based data classification extending the work of Barni et al. in terms of privacy.

DeepSecure [45] is another work in the 2-PC setting, and the passive security model, using Garbled-Circuit techniques. A direct performance comparison of DeepSecure versus CryptoNets [20] confirmed a significant efficiency improvement achieved by DeepSecure.

The recently proposed MiniONN [34] is one of the latest NN-based data classification approaches in the 2-PC setting. MiniONN demonstrates significant performance increase compared to CryptoNets, without loss of accuracy; as well as better accuracy compared to SecureML [35], combined with increased performance. However, it still operates in the 2-PC setting, which is more restricted than the MPC setting we consider, and it only offers passive security.

Under a comparable configuration as MiniONN, and still in the passive security model, Chameleon [43] achieves a 4.2 times performance improvement. Chameleon operates in the 2-PC setting, under the assumption that a Semi-Honest Third Party (STP) is engaged in the offline phase to generate correlated randomness. Despite the strong STP assumption, Chameleon does not need the third party to be involved in the online phase, while it gets a significant performance increase by deploying this STP.

Gazelle [26], the latest work on secure NN classification, outperforms, in terms of efficiency, the best solutions in the literature [20, 34, 43], by carefully selecting which parts of the CNN to carry out using a packed additively homomorphic encryption, and which using garbled circuits. EPIC performs better than Gazelle, while also being secure in the active security model. This is because EPIC only treats linear computations in the privacy-preserving domain, by avoiding NN classification.

To the best of our knowledge, we are the first to provide a privacy-preserving image classification tool combining SVM classification with transfer learning feature extraction, offering active security. EPIC is more efficient than previous work and achieves prediction accuracy higher than that of the related work on the same datasets, although it does not deploy sophisticated NN-based classification on the private inputs. Interestingly, EPIC is not limited to the 2-PC setting, allowing a broad range of application scenarios to be treated by our solution.

3 EPIC

The proposed private image classification solution, EPIC, is based on transfer learning techniques [49] for the feature extraction. As explained in the introduction our main (first) variant of EPIC deploys TensorFlow [1]. In particular, we use the Inception-v3 [47] CNN for feature extraction. We also present a second variant of EPIC, which aims at allowing a trade-off between the accuracy of the classifier’s predictions, and its performance. It does so by deploying a kernel approximation method, on top of Inception-v3 features for dimensionality reduction.

In the second variant of EPIC the algorithm provider needs to publish the feature size of a point and also a parameter γ . At first sight it might seem that γ reveals information about the training data, but we noticed that for our datasets one can fix γ to a small value and modify the regularization parameter C of the SVM. This parameter C will always remain private to the algorithm provider, hence there is no information leakage. We stress again that for both cases, the CNN feature extraction is input independent, so privacy is maintained for the image holder and algorithm provider. Either variant of EPIC prepares the inputs for the next stage, which is the deployment of the Support Vector Machine (SVM). We reiterate that our method assumes that the training of the classifier is performed on cleartext data, prior to our protocol execution.

Specifically, the classification process, depicted in Figure 1, starts with the image holder, who has an image to be classified. The first step consists of the Inception-v3 feature extraction [47] by taking the next to last layer, which has a feature size of 2048, and it is performed by the image holder locally on the cleartext version of the image. For the second variant of EPIC, the kernel approximation sub-step is also considered to be part of the feature extraction phase. The resulting features are then shared (via the secure gateway) to the MPC servers by the image holder, and thus are kept secret from the classification algorithm provider. We indicate secret shared (and thus protected) data in square brackets (Figure 1).

The classification algorithm provider has already trained their SVM classifier. The necessary parameters for the SVM classification, which is the subsequent step, are shared to the MPC servers by the classification algorithm provider and are never revealed to the image holder (nor the analyst). In the following we detail the aforementioned steps, while emphasizing which operations are performed securely.

Note that although EPIC does not allow any information leakage about the private SVM parameters, recent work by Tramer et al. [50] showed that only black-box access to the classifiers can still serve to recover an (near-)equivalent model. We consider this problem to be out of this work’s scope, as it can easily be tackled by restricting the number of queries an external party is allowed to perform on the MPC

Engine. This type of attacks has not been averted by any of the privacy-preserving solutions in the related work.

3.1 Inception-v3 Feature Extraction

High quality features are key to the accuracy of a trained classifier. We ensure high quality feature extraction by deploying the techniques of transfer learning. Specifically, we perform feature extraction based on Inception-v3 [47], which is a *public* CNN classifier trained on a set of non-privacy-sensitive object recognition tasks. Commonly, the training is performed on large datasets, which enhances the prediction accuracy of the classifier. In our context, the trained classifier extracts features based on the activation of a deep convolutional network. Our work shows that powerful feature extraction is essential to the quality of the final classification accuracy. In fact, we demonstrate that the high-complexity (CNN) tasks can be learned on non-private datasets, and still use their power for feature extraction of unrelated tasks. Eventually, this allows us to deploy only linear functions for the actual classification, which enables accurate, and efficient privacy-preserving solutions.

3.2 SVM Classification

SVM classification is one of the most popular classification methods in computer vision. Despite the increasing popularity and high effectiveness of CNN classification techniques, their deployment requires large training datasets [14] that are potentially difficult to obtain when the underlying data is privacy sensitive. Yet, with the CNN features, an SVM can learn quickly from very few positive examples, which shows that they are useful to perform one-shot learning [15].

In addition, black-box transformation of these methods to their privacy-preserving equivalents will result in classifiers that are computationally prohibitive to use. Thus, we opted for the design of a private SVM classifier, while using the techniques of CNN-based transfer learning in the context of feature extraction, which does not raise privacy concerns.

To classify a new unlabeled input with our classifier trained with a linear SVM, we need to securely evaluate the following equation:

$$\text{class}(\mathbf{h}) = \arg \max_i (\mathbf{x}_i \cdot \mathbf{h} + b_i), \quad (1)$$

where:

- \mathbf{h} is the vector representing the client’s image, and has been provided to the MPC servers in shared form;
- b_i is the model intercept (also known as bias), calculated by the classification algorithm provider during the learning phase and secret shared to the MPC servers;
- \mathbf{x}_i are the n support vectors.

The support vectors \mathbf{x}_i , as well as the model intercepts b_i should be protected, as part of the intellectual property of the learned model, and are therefore available to the MPC

EPIC Protocol with kernel approximation as feature reduction

Setup:

1. Algorithm Provider (AP) broadcasts the type of CNN used for feature extraction.
2. AP computes γ from Fig.3 on its own training data. Then AP broadcasts the Init variables from Fig. 3 and secret shares the support vectors x_i, b_i to the MPC engine (Eng). These are stored on Eng as $\llbracket x_i \rrbracket, \llbracket b_i \rrbracket$.

Evaluate:

1. Image Holder (IH) uses public CNN to extract features h' from its image. Then IH maps $h' \mapsto h$ locally using the RBF sampler initialized with the γ broadcasted by AP to obtain a smaller number of features. The new point h is further secret shared to the Eng and stored as $\llbracket h \rrbracket$.
2. Eng uses $\llbracket x_i \rrbracket, \llbracket b_i \rrbracket, \llbracket h \rrbracket$ to compute Equation 1 with a shared result: $\llbracket \text{class}(\mathbf{h}) \rrbracket$.

Figure 2. Protocol for SVM classification with RBF sampler

servers in secret shared form. On the other hand, the client wishes to protect the features extracted from their image. Thus, \mathbf{h} is also secret shared to the MPC servers, by the image holder. Having all the necessary values available, the MPC servers securely calculate the classification label for a new input, following Equation 1. Note that unlike previous work [4, 20], we do not treat the classifier parameters in the clear, but in a secret shared form.

3.3 EPIC with Feature Reduction

The second, and most efficient variant of EPIC protocol is summarized in Figure 2. This EPIC variant trades a small percentage of the classification accuracy to increase efficiency. It achieves this tradeoff by deploying the kernel approximation dimensionality reduction that we describe in Section 3.3.1. Specifically, the protocol starts with the Setup phase, where the algorithm provider (AP) performs the kernel approximation (see Fig. 3) on its own dataset, and broadcasts the type of CNN used, and the Init parameters (see Fig. 3) necessary for the feature reduction at the image holder (IH) side. Then, it secret shares the SVM parameters to the MPC Engine (Eng). In the evaluation phase, the IH performs the feature extraction locally (given the previously obtained parameters), and secret shares the new point to be classified to the Eng. Given all the aforementioned values, the MPC protocol takes place at the Eng, to calculate Equation 1.

Note that if the AP wishes to avoid publishing any randomness generated as $\text{state}_{\text{random}}$ (see Fig. 3), then this randomness can be generated by the Eng, and be kept as a public variable to be used both by the AP, and the IH.

3.3.1 Kernel Approximation

To achieve efficient training of kernel machines (such as SVM) aimed at non-linear problems, several approximation methods (e.g., the method of Rahimi and Recht [39]) have

been proposed. Such approaches have the goal to alleviate the (cleartext) computational, and storage cost of the training, incurred by the high dimensionality of the data, especially when the training datasets are large. The approximation generally is implemented by mapping the input data to a low-dimensional feature space, in such a way that the inner products of the mapped data are approximately equal to the features of a more complex (e.g., Gaussian) kernel. This is known as the kernel trick. These features are later on combined with linear techniques (e.g., linear SVM), yielding an efficient training, but also an efficient classification, which we are able to implement in a privacy-preserving way.

One of the first successful approaches for kernel approximations, achieving high accuracy, was proposed by Rahimi and Recht [39], and is based on random features, which are independent of the training data. To the contrary, Nyström based kernel approximations [16, 53], are data dependent. Although Nyström approximations outperform randomly extracted features [55] in terms of accuracy, the fact that they are data dependent makes them unfit for our purposes, as they can only be computed by applying non-linear functions on the private inputs. From a computational, and storage efficiency perspective, data independent approximations are favored.

We discovered that a variant of the method proposed by Rahimi and Recht [40] is presented in the scikit-learn package [37]. This implements an RBF (Radial Basis Function) sampler, which allows to transform the features without using the training data. This type of dimensionality reduction (much like the feature extraction) is deployed both for the training, and for the data classification. Since the feature selection is random (i.e., data independent), it can be performed on the cleartext data, both by the classification algorithm provider, and by the client, without raising privacy concerns. The second variant of EPIC supports dimensionality reduction for free, by placing all the computational load on the cleartext.

The algorithm implementing the RBF sampler is listed in Figure 3. In our application scenario, the algorithm provider broadcasts the RBF sampler parameters, namely the γ parameter and the feature size. The γ parameter does not reveal any information about the dataset, and we actually used the same value of $\gamma = 2^{-13}$ for all datasets. Note that in Figure 3, γ is a floating point number, which is varied to match a cross-validation score on the training data. The shape variable is the feature size of a point –in our case shape is always going to be 2048– which is the output of Inception-v3 CNN.

4 Experiments

We evaluate EPIC both in terms of accuracy and in terms of performance (i.e., execution time), on three typical benchmark datasets.

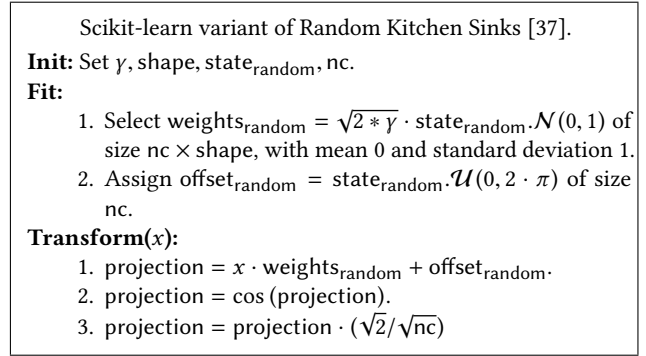


Figure 3. RBF Sampler

4.1 Experimental Setup

Our experiments are conducted on two MPC servers, which yields the most efficient solution, but we also show how the proposed system scales with more than two MPC servers, in Section 4.7. We assume a protocol-independent, input-independent preprocessing phase that takes place prior to the actual protocol execution between the MPC servers. The inputting parties do not need to be aware, nor contribute to this phase. This preprocessing phase creates the necessary randomness to boost the efficiency of the online phase, and allows the inputting parties (image holder and classification algorithm provider) to securely share their inputs.

The online phase begins with the image holder, and the algorithm provider sharing their inputs (reduced CNN features, and SVM parameters, respectively) to the MPC servers. This is performed by executing an interactive protocol between each inputting party and the two MPC servers, as Damgård et al. [11] proposed. Then, the actual private image classification task is executed only between the two MPC servers, as in Figure 2 **Evaluate** phase. In the end, each MPC server sends their resulting share to the image holder, or the analyst, who can then combine the shares and reconstruct the cleartext result, which is the desired class label.

4.2 From Fixed Point Arithmetic to Integers

Looking carefully at the norms, we reduced the underlying field size from 128-bit size prime to 64-bit prime. For the secure comparison sub-protocols that EPIC deploys, we selected the statistical security parameter to be 40 bits. We stress that everywhere the computational security parameter is set to 128. We observed experimentally after running the scikit-learn’s RBF (see Fig. 3) on top of Inception-v3 that each feature is bounded by $\text{abs}(x_i) \leq 15$ where $\text{len}(x) \leq 2048$.

To avoid the costly fixed point arithmetic, we scale each feature x_i by a factor f , and then perform arithmetic on integers. Particularly, we compute $x_i \cdot f$ and then round it to the nearest lower integer. We varied f , and evaluated the SVM’s accuracy. We experimentally concluded at setting $f = 2^8$, as we observed that for $f > 2^8$ we obtained constant SVM classification accuracy, as if we were working on floating

point numbers, while lowering the scale factor f decreased the accuracy by more than 1%. If $f = 2^8$ then to compute a class score from Equation 1 becomes: $s = \sum_{j=1}^{2048} (2^8 \cdot x_{ij} \cdot 2^8 \cdot h_j) + 2^{16} \cdot b_i$ since we need to scale both support vectors x_i as well features h . Using the fact that each component is bounded by 15 then clearly $s \leq 2^{35}$.

Unfortunately, if our inputs are of 35 bit size then there is no room left to perform the secure comparisons in $\arg \max$ with 40 bits statistical security, since our goal is to make all computations inside \mathbb{F}_p where $p \approx 2^{64}$. Surprisingly, for all the datasets we experimented with, s was less than 20 bits long, because some of the x_{ij} 's are negative. Hence, we could run everything on 64-bit primes with 40-bit statistical security, while ensuring there is no information leak from the comparisons. We can achieve even tighter bounding by normalizing the features using the L2-Norm, after the RBF-Sampler invocation. In our setting this is not necessary, since the expected bound on s is already low (20 bits).

4.3 Feature reduction: RBF or PCA?

From the algorithm provider's point of view the goal is to keep the training data private, so we choose to make public only the parameters of the RBF-Init (Fig. 3). As mentioned in Sec. 3.3.1, publishing γ leaks no information about the training data since if one fixes γ in advance, they can vary the regularization parameter C to modify the accuracy of the linear SVM. In our case we always kept $\gamma = 2^{-13}$, and then just tuned C accordingly.

Our experiments showed that the RBF technique reduces the prediction accuracy, which was expected as it decreases the number of components remaining after reduction (see Table 4). On the other hand, PCA requires to compute the equation: $z = A^T \cdot (h - m_h)$ which takes as input a point h and outputs a point z with a length equal to the number of columns in A (denote this as nc). Note that A, h, m_h are all secret shared across the MPC engine. If we reduce the number of components with PCA to $nc = 256$, then we observed that we get around 5% better accuracy than using RBF with the same feature reduction parameter nc .

The main caveat of PCA is that we need to perform (shape- nc) secure multiplications to convert our features, and then another ($nc \cdot n$) secure multiplications for computing the SVM probabilities. Hence, PCA incurs a total of $nc(\text{shape} + n)$ multiplications in the privacy-preserving domain, whereas RBF just ($nc \cdot n$). In application scenarios where the performance is of higher importance, and classification accuracy can be sacrificed (up to a small percentage $< 8\%$), we can see that our RBF approach is still slightly more accurate than other private CNN's (Table 4: CIFAR-10 accuracy, versus Table 2: accuracy of MiniONN [34], and Gazelle [26]).

4.4 Datasets

We selected three image datasets: CIFAR-10, MIT-67, and Caltech-101, to show how EPIC scales in terms of performance, when increasing the number of classes, and to illustrate its classification accuracy.

CIFAR-10 [29]. This is a dataset of 60000 32x32 color images, out of which 50000 are training images and 10000 are test images. CIFAR-10 features 10 classes of objects, with 6000 images per class.

MIT-67 [38]. MIT-67 has 15620 indoor images from 67 scene categories. We used 80 images per class for training, and the rest of the pictures for testing, since this is the default standard. The accuracy metric used here is the mAP (mean Accuracy Precision), which consists of calculating the average over the accuracies of each class.

Caltech-101 [17]. This dataset contains pictures of objects of 102 categories. Each class has at least 31 images and we chose to use 30 images from each class for the training. The accuracy metric is mAP, just as in MIT-67.

4.5 Training

We performed the training on the cleartext datasets described in Section 4.4. Feature extraction was done after resizing each image to 256x256. We trained Linear SVMs based on the one-versus-all strategy (OvA) [52], because it is more efficient to evaluate n classifiers in MPC instead of $n(n-1)/2$. Note that we chose to avoid the data augmentation trick, and adopted the training method presented in DeCAF [15] using the original datasets, and raw features from Inception-v3 [15]. To find parameters that yield high classification accuracy, we have done a grid search for the γ required in the RBF, and the parameter C , which denotes the size-margin hyperplane for the SVM decision function.

We stress that EPIC achieves a sufficient classification accuracy. Given that EPIC workings have been purposely kept simple enough to allow for efficient secure computations, we consider the classification accuracy of EPIC comparable to that achieved by the state-of-the-art (non-privacy-preserving) works in the ML community. The best classification accuracy to date on the CIFAR-10 dataset is 97.14% [18], while EPIC achieves 88.8%. On the MIT-67 dataset, EPIC achieves 72.2% accuracy, while the state-of-the-art (non-secure) solution [28] reports an accuracy of 83.1%. More interestingly, on the Caltech-101 dataset, to the best of our knowledge, the state-of-the-art accuracy is still 93.42% [23], while EPIC achieves 91.4%.

4.6 Classification Accuracy and Performance Evaluation

We executed our experiments, simulating the two MPC servers on two identical desktop computers equipped with Intel i7-4790 processor, at 3.60 GHz over a 1Gbps LAN with an average round-trip ping of 0.3ms.

Our algorithm, hand matches the one listed in Figure 2, where the **Evaluate** step from Figure 2 was implemented using the SPDZ software [7]. The pre-processing phase for this step was estimated using the LowGear protocol by Keller et al. [27], which is the fastest known protocol to produce triples for multiple parties with active security. We do not report on the timings for the feature extraction and reduction phases, since they can be done in the clear, locally by the external parties, which provide inputs to the MPC engine, and they are not privacy-sensitive tasks.

Tune for accuracy. We evaluated the computational performance, data sent over the network, and classification accuracy of EPIC on the default 2048 length feature from the output of Inception-v3. We report these experiment results in Table 3. It is clear that increasing the number of classes n (from 10, to 67, to 102) has a worsening effect on the performance, as the amount of data sent over the network scales linearly with n . The runtime of the online phase is affected less as n increases. For example, going from 10 classes (CIFAR-10) with 0.005 seconds runtime, to 102 classes (Caltech) with 0.03 seconds runtime, the increasing factor is six, whereas for all other metrics the increasing factor is roughly ten (i.e., linear in the number of classes).

In Table 2 we show that EPIC improves over Gazelle [26] in terms of every relevant metric on CIFAR-10: accuracy with 7%, total communication by 50x, and total runtime by 34x. This is because we start with secret shared (powerful) features obtained from public CNNs, whereas Gazelle [26] starts with an encrypted image. We expect Gazelle’s timings to considerably improve, if they adopt our approach, starting from encrypted features produced by a public CNN.

Tune for performance. To increase the performance of EPIC even further, we tried to minimize the feature size used, while still matching the classification accuracy achieved by Gazelle [26] or MiniONN [34] accuracy for CIFAR-10. In the end, we settled with $nc = 128$, and then performed a grid search on γ for the MIT and Caltech datasets. Our results are reported in Table 4. Since the number of features decreases considerably from 2048 to 128 the timings decrease as well. For example, if we look at the online runtime compared to Gazelle [26], our solution improves by a factor of 700x and the total communication cost decreases by almost 500x. We do recognize that our setting is different from the one considered by Gazelle [26], but we see more the similarities, since the end goal is the same, namely to classify secret shared (or encrypted) images.

Our results indicate that general image recognition, and user’s privacy can go well together. In fact we showed that securing the private classification comes nearly for free. A stronger case for CNN features with a Linear SVM should be considered, as a baseline benchmark is done by Razavian et al. [42].

Other optimizations. Note that one of the major improvements came from running the dot products on multiple threads, and doing the argmax operation in a tree-wise manner to decrease the number of communication rounds required by SPDZ.

4.7 Multiparty setting

We benchmarked EPIC on different number of computers with the RBF-128 variant on the CIFAR-10 dataset and measured throughput (operations per second) for the online and offline phases in Figure 4. For the two party case our protocol can carry around 2650 evaluations per second. The throughput decreases with a growing number of parties and reaches 870 ops per second for the five parties case. Notice that the main bottleneck when executing these protocols is still the preprocessing phase, generating the necessary triples.




Figure 4. Throughput of CIFAR-10 evaluations of secret features with RBF-128 EPIC for multiple parties.

4.8 Similar work

It is worth mentioning that we also implemented the method of Barnett et al. [4] in SPDZ, after fixing some security bugs such as cleartext PCA coefficients. They report 124s for one binary classification thus to extrapolate this to 10 classes takes roughly 1240s. To translate the work for Barnett et al. in SPDZ we used a feature extraction algorithm based on Histogram of Oriented Gradients (HOG) and then reduced their dimension using PCA. The reduced points were then plugged into a polynomial SVM to classify the inputs. This methodology yielded a 6.7s execution time of the online phase, and an expensive pre-processing phase of 12 hours for CIFAR-10. The classification accuracy was also poor, namely 58%. This showed that the input dependent phase in MPC is faster than the FHE one, by at least two orders of magnitude, and confirms that our EPIC solution outperforms traditional attempts at classifying images using SVMs.

Framework	Runtime (s)			Communication (MB)			Accuracy %
	Offline	Online	Total	Offline	Online	Total	
MiniONN [34]	472	72	544	3046	6226	9272	81.61
Gazelle [26]	9.34	3.56	12.9	940	296	1236	81.61
EPIC	0.36	0.005	0.37	24	0.33	24.33	88.8

Table 2. 1 Gbps LAN timings for CIFAR-10 dataset on different frameworks.

Dataset	Runtime (s)			Communication (MB)			Accuracy %
	Offline	Online	Total	Offline	Online	Total	
CIFAR-10	0.36	0.005	0.37	24	0.33	24.33	88.8
MIT-67	2.43	0.02	2.45	161.94	2.24	164.18	72.2
Caltech-101	3.71	0.03	3.74	246.59	3.41	250	91.4

Table 3. 1 Gbps LAN timings for EPIC on different datasets with a LinearSVM and normalized features

Dataset	Runtime (s)			Communication (MB)			Accuracy %
	Offline	Online	Total	Offline	Online	Total	
CIFAR-10	0.037	0.0003	0.037	2.472	0.027	2.5	81.74
MIT-67	0.259	0.002	0.261	17.22	0.180	17.4	64.4
Caltech-101	0.395	0.004	0.399	26.27	0.273	26.543	85.56

Table 4. 1 Gbps LAN timings for EPIC on different datasets with a RBF-SVM and a 128 feature size

5 Conclusion and Future Work

We have introduced EPIC, a private image classification system, trained with SVM, while having the input features extracted based on the techniques of transfer learning. We showed how to achieve privacy-preserving image classification in such a way that the task can be fully outsourced to a third, independent party. For our solution we deployed generic MPC tools and showed how to avoid the restricted two-party setting. Unlike all previous work, our approach provides active security, does not leak any information about the private images, nor about the classifier parameters, and it is orders of magnitude more efficient than the privacy-preserving classification solutions proposed in the literature.

Due to their highly accurate predictions, especially for multiclass classification tasks, CNNs have superseded SVM as the state-of-the-art for image classification. However, our work shows that in the privacy-preserving domain, SVM classification can still produce accurate results, as long as it is provided with high quality features. Thus, we have chosen to focus on improving the feature extraction phase, using a transfer learning, CNN-based approach, while avoiding the execution of such complex functions in the MPC domain. An interesting advantage of our solution is that it can be easily applied to the homomorphic encryption domain, since

performing the linear operations has depth 1, and the costlier operation is computing the argmax, which requires to branch on secret comparisons.

Our experiments confirmed that there is a trade-off between the complexity, and therefore also accuracy of the classification algorithms used, versus the efficiency of the privacy-preserving variants of the proposed solutions. In the active security model that we consider in this work, deploying CNNs in the same manner as they are used on cleartext data, is computationally prohibitive with current privacy-preserving methods.

We avoided experimenting with the MNIST dataset, because it consists of gray images, and does not correspond to the distribution of the ImageNet challenge images. The accuracy can improve if we retrain the last layer of Inception-v3. We consider interesting to investigate how to securely retrain some of the CNN nodes in order to fine-tune the network for different applications, such as medical imaging, targeted surveillance, detect fakes or identify faulty models like cars and other devices. To this end, Gazelle [26] is orthogonal to our work, since their solution can also be used as a building block for the evaluation of the last retrained (and secret) layer. There is also some work done by Dahl [10], who points out that evaluating the last layer with secret inputs is

possible. However, we consider this to require more careful investigation and further optimizations to achieve practical timings.

Acknowledgments

This material is based upon work supported by the Defense Advanced Research Projects Agency (DARPA) and Space and Naval Warfare Systems Center, Pacific (SSC Pacific) under contract No. N66001-15-C-4070.

References

- [1] Martín Abadi, Paul Barham, Jianmin Chen, Zhifeng Chen, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Geoffrey Irving, Michael Isard, et al. 2016. TensorFlow: A System for Large-Scale Machine Learning. In *OSDI*, Vol. 16. 265–283.
- [2] Louis JM Aslett, Pedro M Esperança, and Chris C Holmes. 2015. Encrypted Statistical Machine Learning: New Privacy Preserving Methods. *arXiv preprint arXiv:1508.06845* (2015).
- [3] Hossein Azizpour, Ali Sharif Razavian, Josephine Sullivan, Atsuto Maki, and Stefan Carlsson. 2016. Factors of Transferability for a Generic Convnet Representation. *IEEE transactions on pattern analysis and machine intelligence* 38, 9 (2016), 1790–1802.
- [4] Anthony Barnett, Jay Santokhi, Michael Simpson, Nigel P Smart, Charlie Stainton-Bygrave, Srnivas Vivek, and Adrian Waller. 2017. Image Classification using non-linear Support Vector Machines on Encrypted Data. *IACR Cryptology ePrint Archive 2017* (2017), 857. <http://eprint.iacr.org/2017/857>
- [5] Mauro Barni, Claudio Orlandi, and Alessandro Piva. 2006. A Privacy-Preserving Protocol for Neural-Network-Based Computation. In *Proceedings of the 8th workshop on Multimedia and security*. ACM, 146–151.
- [6] Raphael Bost, Raluca Ada Popa, Stephen Tu, and Shafi Goldwasser. 2015. Machine Learning Classification over Encrypted Data. In *NDSS*.
- [7] Bristol Crypto. 2018. SPDZ-2: Multiparty computation with SPDZ, MASCOT, and Overdrive offline phases. <https://github.com/bristolcrypto/SPDZ-2>. (2018).
- [8] Hervé Chabanne, Amaury de Wargny, Jonathan Milgram, Constance Morel, and Emmanuel Prouff. 2017. Privacy-Preserving Classification on Deep Neural Network. *IACR Cryptology ePrint Archive 2017* (2017), 35.
- [9] Melissa Chase, Ran Gilad-Bachrach, Kim Laine, Kristin Lauter, and Peter Rindal. 2017. Private Collaborative Neural Network Learning. *IACR Cryptology ePrint Archive 2017* (2017), 762. <http://eprint.iacr.org/2017/762>
- [10] Morten Dahl. 2018. Private Image Analysis with MPC: Training CNNs on Sensitive Data using SPDZ. <https://mortendahl.github.io/2017/09/19/private-image-analysis-with-mpc/>. (2018).
- [11] Ivan Damgård, Kasper Damgård, Kurt Nielsen, Peter Sebastian Nordholt, and Tomas Toft. 2015. Confidential Benchmarking based on Multiparty Computation. *IACR Cryptology ePrint Archive 2015* (2015), 1006.
- [12] Ivan Damgård, Marcel Keller, Enrique Larraia, Valerio Pastro, Peter Scholl, and Nigel P Smart. 2013. Practical Covertly Secure MPC for Dishonest Majority—or: Breaking the SPDZ Limits. In *European Symposium on Research in Computer Security*. Springer, 1–18.
- [13] Ivan Damgård, Valerio Pastro, Nigel Smart, and Sarah Zakarias. 2012. Multiparty Computation from Somewhat Homomorphic Encryption. In *Advances in Cryptology—CRYPTO 2012*. Springer, 643–662.
- [14] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. 2009. Imagenet: A Large-Scale Hierarchical Image Database. In *Computer Vision and Pattern Recognition, 2009. CVPR 2009. IEEE Conference on. IEEE*, 248–255.
- [15] Jeff Donahue, Yangqing Jia, Oriol Vinyals, Judy Hoffman, Ning Zhang, Eric Tzeng, and Trevor Darrell. 2014. DeCAF: A Deep Convolutional Activation Feature for Generic Visual Recognition. In *International conference on machine learning*. 647–655.
- [16] Petros Drineas and Michael W Mahoney. 2005. On the Nyström Method for Approximating a Gram Matrix for Improved Kernel-Based Learning. *journal of machine learning research* 6, Dec (2005), 2153–2175.
- [17] Li Fei-Fei, R Fergus, and P Perona. [n. d.]. Learning Generative Visual Models from Few Training Examples: An Incremental Bayesian Approach Tested on 101 Object Categories. In *Computer Vision and Pattern Recognition Workshop, 2004. CVPRW'04. Conference on. IEEE*, 178–178.
- [18] Xavier Gastaldi. 2017. Shake-Shake Regularization. *arXiv preprint arXiv:1705.07485* (2017).
- [19] Craig Gentry. 2009. Fully Homomorphic Encryption using Ideal Lattices. In *STOC*, Vol. 9. 169–178.
- [20] Ran Gilad-Bachrach, Nathan Dowlin, Kim Laine, Kristin Lauter, Michael Naehrig, and John Wernsing. 2016. CryptoNets: Applying Neural Networks to Encrypted Data with high Throughput and Accuracy. In *International Conference on Machine Learning*. 201–210.
- [21] Oded Goldreich, Silvio Micali, and Avi Wigderson. 1987. How to Play any Mental Game. In *Proceedings of the nineteenth annual ACM symposium on Theory of computing*. ACM, 218–229.
- [22] Thore Graepel, Kristin Lauter, and Michael Naehrig. 2012. ML Confidential: Machine Learning on Encrypted Data. In *International Conference on Information Security and Cryptology*. Springer, 1–21.
- [23] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2014. Spatial Pyramid Pooling in Deep Convolutional Networks for Visual recognition. In *European conference on computer vision*. Springer, 346–361.
- [24] Gao Huang, Zhuang Liu, Kilian Q Weinberger, and Laurens van der Maaten. 2016. Densely Connected Convolutional Networks. *arXiv preprint arXiv:1608.06993* (2016).
- [25] Yangqing Jia, Evan Shelhamer, Jeff Donahue, Sergey Karayev, Jonathan Long, Ross Girshick, Sergio Guadarrama, and Trevor Darrell. 2014. Caffe: Convolutional Architecture for Fast Feature Embedding. In *Proceedings of the 22nd ACM international conference on Multimedia*. ACM, 675–678.
- [26] Chiraag Juvekar, Vinod Vaikuntanathan, and Anantha Chandrakasan. 2018. GAZELLE: A Low Latency Framework for Secure Neural Network Inference. *arXiv preprint arXiv:1801.05507* (2018).
- [27] Marcel Keller, Valerio Pastro, and Dragos Rotaru. 2018. Overdrive: Making SPDZ great again. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*. Springer, 158–189.
- [28] Fahad Shahbaz Khan, Joost van de Weijer, Rao Muhammad Anwer, Andrew D Bagdanov, Michael Felsberg, and Jorma Laaksonen. 2018. Scale Coding Bag of Deep Features for Human Attribute and Action Recognition. *Machine Vision and Applications* 29, 1 (2018), 55–71.
- [29] Alex Krizhevsky and Geoffrey Hinton. 2009. Learning Multiple Layers of Features from Tiny Images. (2009).
- [30] Sven Laur, Helger Lipmaa, and Taneli Mielikäinen. 2006. Cryptographically Private Support Vector Machines. In *Proceedings of the 12th ACM SIGKDD international conference on Knowledge discovery and data mining*. ACM, 618–624.
- [31] Keng-Pei Lin and Ming-Syan Chen. 2010. Privacy-Preserving Outsourcing Support Vector Machines with Random Transformation. In *Proceedings of the 16th ACM SIGKDD international conference on Knowledge discovery and data mining*. ACM, 363–372.
- [32] Keng-Pei Lin and Ming-Syan Chen. 2011. On the Design and Analysis of the Privacy-Preserving SVM Classifier. *IEEE Transactions on Knowledge and Data Engineering* 23, 11 (2011), 1704–1717.
- [33] Yehuda Lindell and Benny Pinkas. 2000. Privacy Preserving Data Mining. In *Advances in Cryptology—CRYPTO 2000*. Springer, 36–54.

- [34] Jian Liu, Mika Juuti, Yao Lu, and N Asokan. 2017. Oblivious Neural Network Predictions via MiniONN Transformations. In *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*. ACM, 619–631.
- [35] Payman Mohassel and Yupeng Zhang. 2017. SecureML: A System for Scalable Privacy-Preserving Machine Learning. *IACR Cryptology ePrint Archive 2017* (2017), 396.
- [36] Claudio Orlandi, Alessandro Piva, and Mauro Barni. 2007. Oblivious Neural Network Computing via Homomorphic Encryption. *EURASIP Journal on Information Security* 2007 (2007), 18.
- [37] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. 2011. Scikit-learn: Machine Learning in Python. *Journal of Machine Learning Research* 12 (2011), 2825–2830.
- [38] Ariadna Quattoni and Antonio Torralba. 2009. Recognizing Indoor Scenes. In *Computer Vision and Pattern Recognition, 2009. CVPR 2009. IEEE Conference on*. IEEE, 413–420.
- [39] Ali Rahimi and Benjamin Recht. 2008. Random Features for Large-Scale Kernel Machines. In *Advances in neural information processing systems*. 1177–1184.
- [40] Ali Rahimi and Benjamin Recht. 2009. Weighted Sums of Random Kitchen Sinks: Replacing Minimization with Randomization in Learning. In *Advances in neural information processing systems*. 1313–1320.
- [41] Yogachandran Rahulamathavan, Raphael C-W Phan, Suresh Veluru, Kanapathippillai Cumanan, and Muttukrishnan Rajarajan. 2014. Privacy-Preserving Multi-Class Support Vector Machine for Outsourcing the Data Classification in Cloud. *IEEE Transactions on Dependable and Secure Computing* 11, 5 (2014), 467–479.
- [42] Ali Sharif Razavian, Hossein Azizpour, Josephine Sullivan, and Stefan Carlsson. 2014. CNN Features off-the-shelf: An Astounding Baseline for Recognition. In *Computer Vision and Pattern Recognition Workshops (CVPRW), 2014 IEEE Conference on*. IEEE, 512–519.
- [43] M Sadegh Riazi, Christian Weinert, Oleksandr Tkachenko, Ebrahim M Songhori, Thomas Schneider, and Farinaz Koushanfar. 2018. Chameleon: A Hybrid Secure Computation Framework for Machine Learning Applications. *arXiv preprint arXiv:1801.03239* (2018).
- [44] Ronald L Rivest, Len Adleman, and Michael L Dertouzos. 1978. On Data Banks and Privacy Homomorphisms. *Foundations of secure computation* 4, 11 (1978), 169–180.
- [45] Bitu Darvish Rouhani, M Sadegh Riazi, and Farinaz Koushanfar. 2017. DeepSecure: Scalable Provably-Secure Deep Learning. *arXiv preprint arXiv:1705.08963* (2017).
- [46] Reza Shokri and Vitaly Shmatikov. 2015. Privacy-Preserving Deep learning. In *Proceedings of the 22nd ACM SIGSAC conference on computer and communications security*. ACM, 1310–1321.
- [47] Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. 2016. Rethinking the Inception Architecture for Computer Vision. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 2818–2826.
- [48] Sin G Teo, Shuguo Han, and Vincent CS Lee. 2013. Privacy Preserving Support Vector Machine using non-linear Kernels on Hadoop Mahout. In *Computational Science and Engineering (CSE), 2013 IEEE 16th International Conference on*. IEEE, 941–948.
- [49] Sebastian Thrun. 1996. Is Learning the n-th Thing any Easier than Learning the First?. In *Advances in neural information processing systems*. 640–646.
- [50] Florian Tramèr, Fan Zhang, Ari Juels, Michael K Reiter, and Thomas Ristenpart. 2016. Stealing Machine Learning Models via Prediction APIs. In *USENIX Security Symposium*. 601–618.
- [51] Jaideep Vaidya, Hwanjo Yu, and Xiaoqian Jiang. 2008. Privacy-Preserving SVM Classification. *Knowledge and Information Systems* 14, 2 (2008), 161–178.
- [52] Vladimir Naumovich Vapnik and Vlamimir Vapnik. 1998. *Statistical Learning Theory*. Vol. 1. Wiley New York.
- [53] Christopher KI Williams and Matthias Seeger. 2001. Using the Nyström Method to Speed up Kernel Machines. In *Advances in neural information processing systems*. 682–688.
- [54] Linjie Yang, Ping Luo, Chen Change Loy, and Xiaoou Tang. 2015. A large-scale car dataset for fine-grained categorization and verification. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 3973–3981.
- [55] Tianbao Yang, Yu-Feng Li, Mehrdad Mahdavi, Rong Jin, and Zhi-Hua Zhou. 2012. Nyström Method vs Random Fourier Features: A Theoretical and Empirical Comparison. In *Advances in neural information processing systems*. 476–484.
- [56] Andrew Chi-Chih Yao. 1986. How to Generate and Exchange Secrets. In *Foundations of Computer Science, 1986., 27th Annual Symposium on*. IEEE, 162–167.
- [57] Jason Yosinski, Jeff Clune, Yoshua Bengio, and Hod Lipson. 2014. How Transferable are Features in Deep Neural Networks?. In *Advances in neural information processing systems*. 3320–3328.