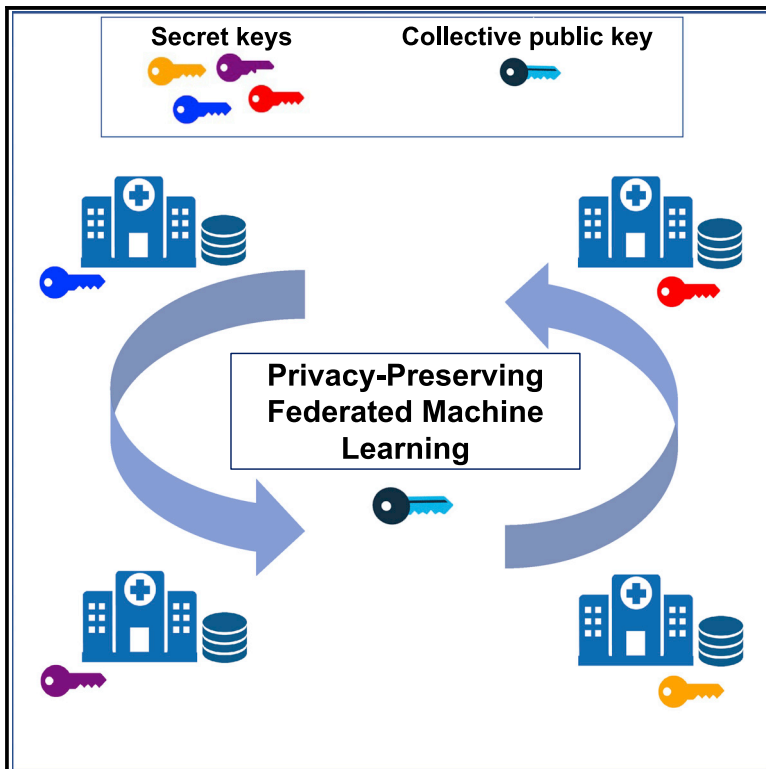


Patterns

Privacy-preserving federated neural network learning for disease-associated cell classification

Graphical abstract



Highlights

- We enable collaborative and privacy-preserving model training between institutions
- Training under encryption does not degrade the utility of the data
- We apply our solution to the single-cell analysis in a federated setting
- Our method is generalizable to other machine learning tasks in the healthcare domain

Authors

Sinem Sav, Jean-Philippe Bossuat,
Juan R. Troncoso-Pastoriza,
Manfred Claassen,
Jean-Pierre Hubaux

Correspondence

sinem.sav@epfl.ch (S.S.),
manfred.claassen@
med.uni-tuebingen.de (M.C.),
jean-pierre.hubaux@epfl.ch (J.-P.H.)

In brief

To enable federated learning with several healthcare institutions in a privacy-preserving way, this work relies on multiparty homomorphic encryption. The paper describes several optimizations for enabling efficient and collaborative machine learning training under encryption. The authors demonstrate the utility and performance of the proposed solution with an application to disease-associated cell classification in a federated setting. The solution preserves the utility while protecting the privacy of the institutions' data, the model, and the exchanged values between the institutions.



Article

Privacy-preserving federated neural network learning for disease-associated cell classification

Sinem Sav,^{1,5,*} Jean-Philippe Bossuat,² Juan R. Troncoso-Pastoriza,² Manfred Claassen,^{3,4,*} and Jean-Pierre Hubaux^{1,2,*}¹Laboratory for Data Security (LDS), EPFL, Lausanne 1015, Switzerland²Tune Insight SA, Lausanne 1015, Switzerland³Internal Medicine I, University Hospital Tübingen, Faculty of Medicine, University of Tübingen, Tübingen 72016, Germany⁴Department of Computer Science, University of Tübingen, Tübingen 72076, Germany⁵Lead contact*Correspondence: sinem.sav@epfl.ch (S.S.), manfred.claassen@med.uni-tuebingen.de (M.C.), jean-pierre.hubaux@epfl.ch (J.-P.H.)<https://doi.org/10.1016/j.patter.2022.100487>

THE BIGGER PICTURE High-quality medical machine learning models will benefit greatly from collaboration between health care institutions. Yet, it is usually difficult to transfer data between these institutions due to strict privacy regulations. In this study, we propose a solution, *PriCell*, that relies on multiparty homomorphic encryption to enable privacy-preserving collaborative machine learning while protecting via encryption the institutions' input data, the model, and any value exchanged between the institutions. We show the maturity of our solution by training a published state-of-the-art convolutional neural network in a decentralized and privacy-preserving manner. We compare the accuracy achieved by *PriCell* with the centralized and non-secure solutions and show that *PriCell* guarantees privacy without reducing the utility of the data. The benefits of *PriCell* constitute an important landmark for real-world applications of collaborative training while preserving privacy.



Development/Pre-production: Data science output has been rolled out/validated across multiple domains/problems

SUMMARY

Training accurate and robust machine learning models requires a large amount of data that is usually scattered across data silos. Sharing or centralizing the data of different healthcare institutions is, however, unfeasible or prohibitively difficult due to privacy regulations. In this work, we address this problem by using a privacy-preserving federated learning-based approach, *PriCell*, for complex models such as convolutional neural networks. *PriCell* relies on multiparty homomorphic encryption and enables the collaborative training of encrypted neural networks with multiple healthcare institutions. We preserve the confidentiality of each institutions' input data, of any intermediate values, and of the trained model parameters. We efficiently replicate the training of a published state-of-the-art convolutional neural network architecture in a decentralized and privacy-preserving manner. Our solution achieves an accuracy comparable with the one obtained with the centralized non-secure solution. *PriCell* guarantees patient privacy and ensures data utility for efficient multi-center studies involving complex healthcare data.

INTRODUCTION

Machine learning models, in particular neural networks, extract valuable insights from data and have achieved unprecedented predictive performance in the healthcare domain, e.g., in single-cell analysis,⁵ aiding medical diagnosis and treatment,^{6,7} or in personalized medicine.⁸ Training accurate and unbiased models without overfitting requires access to a large amount of diverse data that is usually isolated and scattered across

different healthcare institutions.⁹ Sharing or transferring personal healthcare data is, however, often unfeasible or limited due to privacy regulations, such as General Data Protection Regulation (GDPR) or Health Insurance Portability and Accountability Act (HIPAA). Consequently, privacy-preserving collaborative learning solutions play a vital role for researchers, as they enable medical advances without the information about each institution's data being shared or leaked. Collaborative learning solutions play a particularly important role for studies that involve



novel informative, yet not universally established, data modalities, such as high-dimensional single-cell measurements, where the number of examples is typically low at individual study centers and only amounts to critical mass for the successful training of machine learning models across multiple study centers.¹⁰ The ability to satisfy privacy regulations in an efficient and effective manner constitutes a pivotal requirement to carry out translational multi-center studies.

Federated learning (FL) has emerged as a promising distributed learning approach, where the parties keep their raw data on their premises and exchange intermediate model parameters.¹¹ This approach has enabled collaborative learning for several medical applications, and it has been shown that FL performs comparably with centralized training on medical datasets.^{12–14} Recently, the concept of swarm learning (SL) has been proposed; it enables decentralized machine learning for precision medicine. The seminal work of SL¹⁵ is based on edge computing and permissioned blockchains and removes the need for a central server in the FL approach. Despite the advantages of FL and SL for keeping the sensitive data local and for reducing the amount of data transferred/outsourced, the model and the intermediate values exchanged between the parties remain prone to several privacy attacks executed by the other parties or the aggregator (in FL), such as membership inference attacks^{16,17} or reconstructing the parties' inputs.^{18–20} In this work, we provide a solution that further conceals the global machine learning model from the participants by relying on mathematically secure cryptographic techniques to mitigate these inference attacks.

To mitigate or prevent the leakage in the FL setting, several privacy-preserving mechanisms have been proposed. These mechanisms can be classified under three main categories, depending on the strategy they are based on: differential privacy (DP), secure multiparty computation (SMC), and homomorphic encryption (HE).

DP-based solutions aim to perturb the parties' input data or the intermediate model values exchanged throughout the learning. Several studies in the medical domain keep the data on the local premises and use FL with a differential privacy mechanism on the exchanged model parameters.^{21–23} Despite being a pioneering mitigation against privacy attacks, DP-based solutions perturb the model parameters, thus decreasing the utility and making the deployment harder for medical applications, where the accuracy is already constrained by limited data. Quantification of the privacy achieved via DP-based approaches is also very difficult²⁴ and the implementation of DP, especially in medical imaging applications, is not a trivial task.⁹

Another line of research relies on SMC techniques to ensure privacy and to enable collaborative training of machine learning models.^{25–29} SMC techniques rely on secret-sharing the data of the parties and on performing the training on the secret-shared data among multiple computing nodes (usually 2, 3, or 4 nodes). Nevertheless, it is usually hard to deploy these solutions, as they often rely on a trusted third party for the sake of efficiency. Moreover, their scalability with the number of parties is poor due to the large communication overhead.

Finally, several works employ HE to enable secure aggregation or to secure outsourcing of the training to a cloud server.^{30,31} These solutions, however, cannot solve the distributed scenario where parties keep their local data in their premises.

The adoption of each of the aforementioned solutions introduces several privacy, utility, and performance trade-offs that need to be carefully balanced for healthcare applications. To balance these trade-offs, several works employ multiparty homomorphic encryption (MHE).^{32,33} Although the underlying model in these solutions enables privacy-preserving distributed computations and maintains the local data of the parties on their local premises, the functionality of these works is limited to the execution of simple operations, i.e., basic statistics, counting, or linear regression, and the underlying protocols do not support an efficient execution of neural networks in the FL setting.

Recently, Sav et al. proposed a more versatile solution, POSEIDON, for enabling privacy-preserving federated learning for neural networks by relying on MHE³⁴ to mitigate FL inference attacks by keeping the model and intermediate values encrypted. Their solution, however, does not address the implementation and efficient execution of convolutional neural networks, a widely adopted machine learning model to analyze complex data types, such as single-cell data.

We propose *PriCell*, a solution based on MHE to enable the training of a federated convolutional neural network in a privacy-preserving manner, thus preserving the utility of the data for single-cell analysis. To the best of our knowledge, *PriCell* is the first of its kind in the regime of privacy-preserving multi-center single-cell studies under encryption. By bringing privacy-by-design and by preventing the transfer of patients' data to other institutions, our work contributes to single-cell studies and streamlines the slow and demanding process of the reviewing of independent ethics committees for consent forms and study protocols. To mitigate FL attacks, we keep the model and any value that is exchanged between the parties in an encrypted form, and we rely on the threat model and setting proposed in the work of Sav et al.³⁴ (detailed in the [experimental procedures](#)).

By designing new packing strategies and homomorphic matrix operations, we improve the performance of the protocols for encrypted convolutional neural networks that are predominantly used in the healthcare domain.³⁵ To evaluate our system within the framework of single-cell analysis, we train a convolutional neural network (CellCnn), designed by Arvaniti and Claassen,⁴ within our privacy-preserving system for the disease classification task. We also show the feasibility of our solution with several single-cell datasets utilized for cytomegalovirus infection (CMV)¹ and acute myeloid leukemia (AML)² classification, and one dataset for non-inflammatory neurological disease (NIND) and relapsing-remitting multiple sclerosis (RRMS)³ classification. We compare our classification accuracy in a privacy-preserving FL setting with the centralized and non-encrypted baseline. Our solution converges comparably with the training with centralized data, and we improve on the state-of-the-art decentralized secure solution³⁴ in terms of training time. For example, in a setting with 10 parties, we improve POSEIDON's execution time by at least one order of magnitude.

RESULTS

In this section, we introduce the system overview of our solution, present the neural network architecture that is used for our evaluation, and lay out our experimental findings.

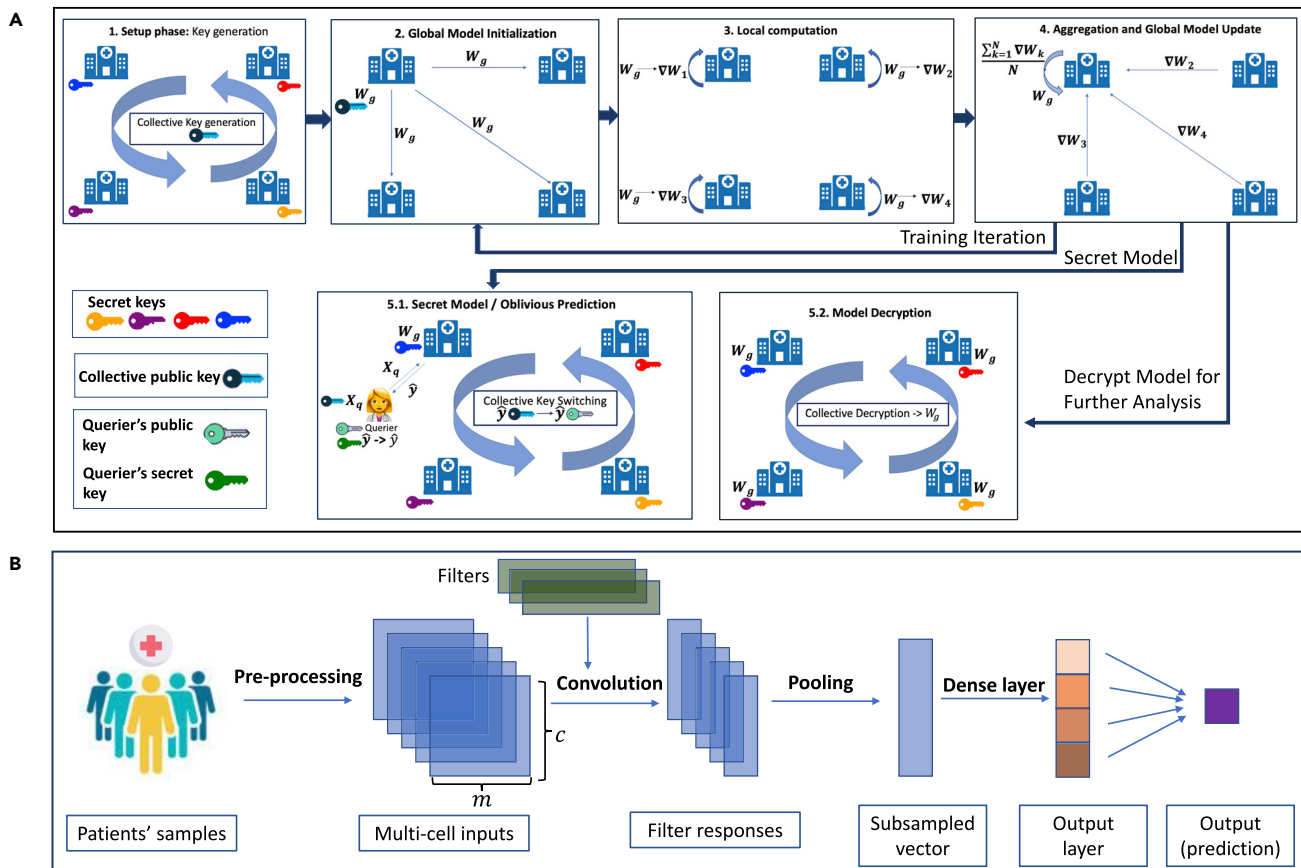


Figure 1. PriCell's system model

(A) PriCell's training and evaluation workflow. Training encapsulates the generation of cryptographic keys and the federated learning iterations on an encrypted model with multiple healthcare institutions. After training, the model is either kept encrypted or decrypted for further analysis.

(B) CellCnn⁴ neural network architecture that is used in the local computation phase of (A). The network takes multi-cell samples as an input and applies a 1D convolution with h filters followed by a pooling layer. A dense (fully connected) layer then outputs the phenotype prediction.

System overview

We summarize PriCell's system and its workflow for collaborative training and query evaluation (prediction) in Figure 1A. Assuming that there are four healthcare institutions with each holding its respective secret key, the workflow starts with the generation of a *collective public key* and a set of *evaluation keys* that are necessary for the encrypted operations, using each participant's secret key. We refer to this phase as the *setup phase*. In the second phase, the participants agree on the initial random global model weights (W_g) and encrypt them with the collective public key. We denote the encryption of any value with boldface letters, i.e., W_g . After encrypting the initial global weights, the *local computation phase* begins; we base this phase on a variant of the FedAvg algorithm¹¹ to enable collective training. Based on this algorithm, to find the model gradients (∇W_k) each party performs several encrypted training iterations on their local data (local iteration). The local model gradients are then sent and aggregated at one of the parties that will perform the global model update. The updated model is then broadcast back and the process returns to phase 2. After a fixed number of training iterations, the participants can choose to keep the model confidential (option 5.1 in

Figure 1A) or to decrypt it for further analysis (option 5.2 in Figure 1A).

If prediction-as-a-service is offered to a querier (a researcher) and the model is kept encrypted, the querier must encrypt the evaluation data (X_q) with the collective public key of the parties. Once the prediction is done, the result (\hat{y}) is collectively switched to the public key of the querier by using the underlying cryptoscheme's collective key-switching functionality. If the model is instead decrypted, the querier encrypts the data with their own key; hence, no key switch is needed after the prediction. As a result, regardless of the model being confidential or not, the evaluation data of the querier and the prediction result always remain protected, as only the querier can decrypt the end result.

CellCnn model overview

CellCnn is a convolutional neural network that enables multi-instance learning and associates a set of observations on cellular population, namely multi-cell inputs, with a phenotype.⁴ This architecture is designed for detecting rare cell subsets associated with a disease by using multi-cell inputs generated from high-dimensional single-cell marker measurements. By their nature, these inputs can be used to predict the phenotype of a donor

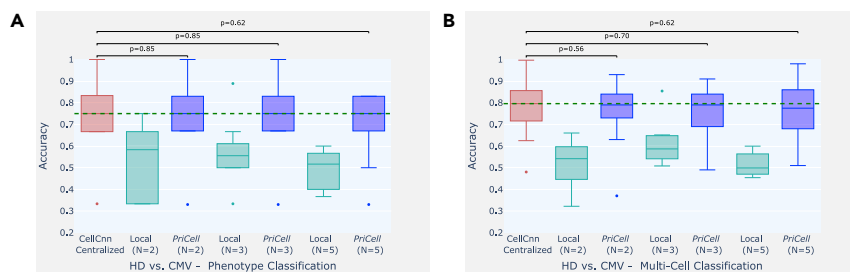


Figure 2. Accuracy boxplots when classifying healthy donor versus cytomegalovirus infection for training multi-cells drawn from the bag of all cells per class

(A and B) Experiments are repeated 10 times with different train and test set splits, the vertical dashed line illustrates the median for the baseline (CellCnn) and the dots represent the outliers. The p values shown at the top of the figure are calculated with a Wilcoxon signed-rank test for the comparison between the corresponding boxplots ($p > 0.05$ indicates that the distributions are not significantly different). Classification accuracy is reported for two datasets: (A) phenotype classification of 6 patients and (B) multi-cell input classification on 4,000 samples. HD, healthy donor; CMV, cytomegalovirus infection.

or the associated phenotype for a given subset of cells. In this scenario, we enable privacy-preserving and distributed multi-instance learning, and we compare our classification performance with the baseline (CellCnn⁴ trained on centralized data with no privacy protection). We note here that replicating the full-pipeline of CellCnn⁴ for downstream analysis requires either heavy approximations under encryption or the decryption of the trained model. Our solution enables the collective and privacy-preserving training for the classification task, whereas subsequent analyses that require access to the model are beyond the scope of this work. Yet, we show the negligible effect that our encryption would practically have on these analyses in Note S6.

We show the architecture of CellCnn⁴ in Figure 1B. The network comprises a 1D convolutional layer followed by a pooling layer and a dense (fully connected) layer. Each multi-cell input sample in Figure 1B is generated using c cells per phenotype with m features (markers), and these samples are batched to construct multi-cell inputs. The training set is then generated by choosing z multi-cell inputs per output label or per patient.

We refer to the reader to the work of Arvaniti and Claassen⁴ for the details of the neural network architecture. We detail the changes we introduce to this architecture to enable operations under HE in the [local neural network operations](#) in the [experimental procedures](#).

Experimental evaluation

We evaluate our proposed solution in terms of model accuracy, runtime performance, scalability with the number of parties, number of data samples, number of features, and communication overhead. In this section, we also provide a comparison with previous work. We give details on the machine learning hyperparameters and security parameters used for our evaluation in Note S3.

Model accuracy

To assess our solution in terms of accuracy, we use the same datasets used in two peer-reviewed biomedical studies.^{4,3} We rely on three datasets to perform NIND, RRMS, CMV, and AML classification. We give the details of each dataset in the [datasets](#) subsection of the [experimental procedures](#). Our aim is to show that PriCell achieves a classification performance on par with the centralized non-private baseline.

As the original studies rely on centralized datasets, we evenly distribute the individual donors in the respective dataset over N

parties. We give the classification performance on these datasets in Figures 2, 3, and 4, and provide a tabular version of all the results corresponding to these plots with accuracy, precision, recall, and F score metrics in Tables S3–S7. The x axis shows different training approaches: (1) the data are centralized and the original CellCnn approach⁴ is used for training and classification to construct a baseline, (2) each party trains a model with only its local data (Local) without collaborating with other parties, and (3) our solution for privacy-preserving collaboration between parties is used (PriCell). For the Local training (2), we average the test accuracy achieved by individual parties. We perform Wilcoxon signed-rank test to prove that the accuracy of the centralized CellCnn and PriCell are not significantly different. $p > 0.05$ indicates that there is not enough evidence to reject the null hypothesis and all our findings suggest that there is no significant difference between PriCell and the centralized non-private baseline (CellCnn). Note that this is the desired outcome as our aim is to show that the results of PriCell and CellCnn are similar.

In our experiments, random multi-cell inputs that are used for training are drawn with replacement from the original training samples. Drawing multi-cell inputs can be done in two ways: using the bag of all cells per class or individually drawing them from each patient. We report the classification performance by using two test datasets: one set is generated using multi-cell inputs with $c = 100 - 200$ cells drawn from all patients in the test set to increase the size of the test set for multi-cell classification, and the second set is generated by drawing 1,000–10,000 cells from each donor separately for phenotype prediction. We give more details about the setting and hyperparameters for each experiment in Note S3.

For CMV classification, we generate the training data by drawing random cell subsets from the cell bags per phenotype. For NIND and RRMS classification, we observe that drawing multi-cells per phenotype varies the accuracy between runs and that the median accuracy over 10 runs increases when distributing the initial dataset among $N = 6$ parties (see Figures 3A, 3B, 4A, and 4B). This suggests that separately drawing multi-cell inputs from each individual performs better for this task, as corroborated by the results obtained with drawing 2,000 cells from each patient with replacement (see Figures 3C, 3D, 4C, and 4D). Finally, in Table S2, we report the median accuracy, precision, recall, and F score of 10 runs (with different train and test set splits) on patient-based sub-sampling for NIND and RRMS,




Figure 3. Accuracy boxplots when classifying healthy donor versus relapsing-remitting multiple sclerosis

(A–D) Accuracy boxplots when classifying healthy donor versus relapsing-remitting multiple sclerosis, for training multi-cells drawn from the bag of all cells per class (A and B) and drawn from each patient separately (C and D).

Experiments were repeated 10 times with different train and test set splits, the vertical dashed line illustrates the median for the baseline (CellCnn) and the dots represent the outliers. The p values shown at the top of the figure are calculated with a Wilcoxon signed-rank test for the comparison between the corresponding boxplots ($p > 0.05$ indicates that the distributions are not significantly different). Classification accuracy is reported for 2 datasets: multi-cell input classification on 96 samples and phenotype classification of 12 patients. HD, healthy donor; RRMS, relapsing-remitting multiple sclerosis.

and phenotype-based sub-sampling for CMV and we provide the F score plots for all experiments in Figures S1–S3 to support our findings.

To construct a realistic overall distribution, we limit the number of parties to be lower than the number of donors in the dataset. We observe that, given a sufficient number of samples per party, our distributed secure solution achieves classification performance comparable with the original work, where the data are centralized and the training is done without privacy protection. In the experiments on CMV, for example, the median accuracy achieved by *PriCell* is exactly the same as the centralized baseline for phenotype classification and very close (at most 2% gap) for multi-cell classification. Analogous results are obtained for the other experiments: our privacy-preserving distributed solution achieves almost the same median accuracy with the baseline in RRMS and NIND with patient-based sub-sampling, where the datasets are sufficiently large to be distributed among up to six parties.

Finally, we provide the classification performance on AML in Table S2. As the dataset is relatively small, emulating a distributed setting with more than two parties was not feasible for this task and, as the accuracy does not vary in between different train-test splits, we do not provide the boxplots on the accuracy. However, we observe that, with two parties in *PriCell* training, the accuracy remains exactly the same as the centralized baseline for AML classification.

Most importantly, our evaluation shows that there is always a significant gain in classification performance when switching from local training to privacy-preserving collaboration. The number of donors that each institution has is insufficient for individually training a robust model. In all experimental settings, for a fixed number of N , *PriCell* achieves better performance than the local training while ensuring the confidentiality of the local data.

Runtime

In Table 1, we report the execution times for the training and prediction with $N = 10$ parties and a ring degree $\mathcal{N} = 2^{15}$. To be able to compare the runtimes at a larger scale, we use synthetically

generated data for this set of experiments and vary the number of features (m). We generate a dataset of 1,000 samples per party with $c = 200$ cells per sample. We use $h = 8$ filters, a local batch size of $n = 100$, and 20 global epochs for training. We report the execution time of the setup phase, of the local computations, and of its communication. We include the execution time of distributed bootstrapping (see the experimental procedures for details) as part of the communication time, which takes 1.2 s per iteration and 122 s over 20 epochs. Hence, the communication column for training comprises the time to perform all communication between parties throughout the training, distributed bootstrapping, and the model update.

We observe that *PriCell* trains, in less than 20 min, a CellCnn model on a training set of 200 cells per sample, 1,000 samples per party, and 32 features across 10 parties, including the setup phase and communication. The training time, when the number of features varies, remains 20–25 min, which is the result of our efficient use of the SIMD (single instruction, multiple data) operations provided by the cryptosystem; this is further discussed in the scalability analysis.

In Table 1, we also report the execution times of an oblivious prediction when both the model and the data are encrypted (phase 5.1 of Figure 1A). We recall that the collective key-switching operation enables us to change the encryption key of a ciphertext from the parties' collective key to the querier's key. The maximum number n of samples that can be batched together for a given ring degree \mathcal{N} , number of labels o , and number of features m , is $(\mathcal{N}/2)/(m \cdot o)$ (we also need $m/2$ ciphertexts to batch those samples; see Note S4 for more details). Hence, in our case, the maximum prediction batch size for $\mathcal{N} = 2^{15}$ and $o = 2$ is $n = 2^{13}/m$.

We observe that the local computation for the prediction increases linearly with m , and is linked to the cost of the dominant operation, the convolution, which is, unlike training, carried out between two encrypted matrices (see Note S4). The communication required for prediction includes $m/2$ ciphertexts sent by the querier and one ciphertext (prediction result) sent back by the server. Hence, the communication time also increases




Figure 4. Accuracy boxplots when classifying healthy donor versus non-inflammatory neurological disease

(A–D) Accuracy boxplots when classifying healthy donor versus non-inflammatory neurological disease for training multi-cells drawn from the bag of all cells per class (A and B) and drawn from each patient separately (C and D).

Experiments were repeated 10 times with different train and test set splits, the vertical dashed line illustrates the median for the baseline (CellCnn) and the dots represent the outliers. The p values shown at the top of the figure are calculated with a Wilcoxon signed-rank test for the comparison between the corresponding boxplots ($p > 0.05$ indicates that the distributions are not significantly different). Classification accuracy is reported for 2 datasets: multi-cell input classification on 96 samples and phenotype classification of 12 patients. HD, healthy donor; NIND, non-inflammatory neurological disease.

linearly with m . Finally, the time for the collective key switch remains constant, as it is performed once at the end of the prediction protocol on only one ciphertext.

Scalability analysis

Figure 5 shows the scalability of *PriCell* with the number of parties, the global number of rows (samples), the number of features (markers), and the number of filters for one global training epoch that is to process once all the data of all parties. Unless otherwise stated, we use $c = 200$ cells per sample, a local batch size of $n = 100$, $m = 38$ features, and $h = 8$ filters, for all settings.

We first report the runtime with an increasing number of parties (N) in Figures 5A and 5B when the global number of data samples is fixed to $s = 18,000$ and when the number of samples per party is fixed to 500, respectively. As the parties perform local computations in parallel, *PriCell*'s runtime decreases with increasing N when s is fixed (Figure 5A). When the number of data samples is constant per party, *PriCell*'s computation time remains almost constant and only the communication overhead increases when increasing N (Figure 5B).

We further analyze *PriCell*'s scalability for $N = 10$ when varying the number of global samples (s), the number of features (m), and the number of filters (h). In Figure 5C, we show that *PriCell* scales linearly when increasing the number of global samples with $N = 10$. Increasing the number of features and filters has almost no effect on *PriCell*'s runtime due to our efficient packing strategy that enables SIMD operations through features and filters. However, we note that the increase in $h = 64$ in Figure 5E is due to increasing the cryptosystem parameter \mathcal{N} to have a sufficient number of slots to still rely on our one-cipher packing strategy. The increase in runtime is still linear with respect to \mathcal{N} and, as expected, the use of larger ciphertexts also produces a slight increase in the communication

Comparison with previous work

The most recent solutions for privacy-preserving FL in the N -party setting use DP, SMC, or HE.

DP-based solutions^{21–23} in the medical domain introduce noise in the intermediate values to mitigate adversarial FL attacks. However, it has been shown that training an accurate

model with DP requires a high-privacy budget.³⁶ Thus, DP-based solutions introduce a privacy-accuracy trade-off by perturbing the model parameters, whereas *PriCell* decouples the accuracy from the privacy and achieves privacy-by-design with a reasonable overhead.

SMC-based solutions^{25–29} often require the data providers to communicate their data outside their premises to a limited number of computing nodes, and these solutions assume an honest majority among the computing nodes to protect the data and/or model confidentiality. Comparatively, *PriCell* scales efficiently with N parties and permits them to keep their data on their premises, withstanding collusions of up to $N - 1$ parties.

Finally, HE-based solutions for privacy-preserving analytics in distributed medical settings^{32,33} allow for functionalities (e.g., basic statistics, counting, or linear regression) different than those that *PriCell* enables, and they do not enable the efficient execution of neural networks in an FL setting. Due to the fact that the underlying system, the threat model, and the enabled functionalities of all the aforementioned solutions are different from *PriCell*, a quantitative comparison with these works is a challenging task.

We build on the system and threat model proposed by Sav et al.³⁴ for enabling privacy-preserving FL for neural networks by relying on MHE and make a quantitative comparison with POSEIDON. *PriCell* improves upon the state-of-the-art solution, POSEIDON, by at least one order of magnitude for training times when the same number of epochs and filters is used. This is due to *PriCell*'s design for optimizing the use of SIMD operations, with a packing strategy that enables encrypting all samples of a batch in a single ciphertext; whereas POSEIDON packs the samples within a batch in different ciphertexts. For a local batch size of $n = 1$, 8 filters, 38 features, and 200 cells per sample, *PriCell*'s local computation time is 1.7 s, whereas POSEIDON's is 15.4 s. Increasing the batch size to 100 results in a 100× slower local execution for POSEIDON, whereas it remains constant for *PriCell*, as all samples are packed in one ciphertext. In summary, increasing the batch size or the number of filters yields a linear increase in the advantage of our solution, in terms of local computation time.

Table 1. *PriCell*'s execution times for training and prediction with a varying number of features (m), 10 parties, and ring degree $N = 2^{15}$ (2^{14} ciphertext slots)

m	Training execution time (s)			Prediction execution time (s)		
	Setup	Local computation	Communication	Local computation querier + server	Communication	Collective key-switch
8	17.8	753.4	370.1	0.2 + 0.1	0.3	0.3
16	18.1	778.7	387.0	0.3 + 0.2	0.6	0.3
32	19.3	836.1	393.7	0.3 + 0.4	1.0	0.3
64	21.9	951.1	373.1	0.6 + 0.6	2.2	0.3
128	24.5	1,135.9	374.8	1.6 + 1.5	4.2	0.3

The computation is single-threaded in a virtual network with an average network delay of 0.17 ms and 1 Gbps bandwidth on 10 Linux servers with an Intel Xeon E5-2680 v.3 CPUs running at 2.5 GHz with 24 threads and 12 cores and 256 GB RAM.

Downstream analysis

The training in the original CellCnn study aims at detecting rare disease-associated cell subsets via further analysis.⁴ Assuming the end model is decrypted upon pre-agreement to conduct these analyses, we further investigate how the changes that we introduce in the CellCnn architecture (see [local neural network operations](#) in the [experimental procedures](#)) affect the detection capability. To be able to make a comparison with the original study in terms of detection capability, we introduce these changes in the original implementation of CellCnn, simulate our encryption, and evaluate the impact in the subsequent analyses. We report our results and how our changes to the circuit and training affect the detection capability on rare CMV infection in [Note S6](#) (see [Figures S4–S6](#)).

DISCUSSION

In this work, we present *PriCell*, a system that enables privacy-preserving federated neural network learning for healthcare institutions, in the framework of an increasingly relevant single-cell analysis, by relying on MHE. To the best of our knowledge, *PriCell* is the first solution to enable the training of convolutional neural networks with N parties under encryption on single-cell data. Using MHE, our solution enables the parties to keep their data on their local premises and to keep the model and any FL intermediate values encrypted end-to-end throughout the training. As such, *PriCell* provides security against inference attacks to FL.^{16–20} *PriCell* also protects the querier's (researcher's) evaluation data by using oblivious prediction. The underlying encryption scheme of *PriCell* provides post-quantum security and does not degrade the utility of the data, contrarily to differential privacy-based solutions.^{21,37}

In this work, we demonstrate the flexibility of *PriCell* with the different learning parameters (e.g., batch size, number of features, number of filters), different real-world datasets, and a varying number of parties. Our empirical evaluation shows that *PriCell* is able to efficiently train a collective neural network with a large number of parties while protecting the model and the data through HE. We also show that *PriCell*'s computation and communication overhead remains either constant or scales linearly with the number of parties and with the model parameters.

Furthermore, we show that *PriCell* achieves classification accuracy comparable with the centralized and non-encrypted training. Our evaluation demonstrates a substantial accuracy

gain by collaboration between the parties when compared with locally training with their data only.

In terms of limitations, *PriCell* relies on the assumption that the parties provide non-tampered data. Although poisoning attacks are beyond the scope of this work, *PriCell* can partially mitigate this kind of threat by integrating several mechanisms, such as statistical checks on parties' input data³⁸ or by integrating zero-knowledge proofs during training.^{39,40} Using such techniques for training under encryption in time-constrained applications remains an open research problem.

While enabling privacy-preserving training over N parties, by default *PriCell* does not monitor the training when the whole process is carried out under encryption. As *PriCell* performs all training operations under encryption, the partial models cannot be assessed. However, relaxing the end-to-end security requirement over training, the parties can collectively decrypt the validation accuracy that is computed on an independent validation set. As for the conditions for *PriCell* to lead to accurate and unbiased models, they remain the same as for standard FL approaches, such as tuning the hyperparameters for client and server. This is a possible future research direction and not a trivial task.⁴¹ While the IID setting is not a prerequisite for *PriCell*, we observe that an increased skewness between the parties' data distribution can decrease the model performance. Similar to poisoning attacks, this limitation can be partially addressed by integrating statistical checks to ensure the skewness is below a certain threshold; this requires a comprehensive study on skew types (i.e., quantity, label, or feature) and the thresholds for each of them for single-cell analysis.

In general, as data sharing in the healthcare domain is usually prevented due to the sensitive nature of data, and due to privacy regulations such as HIPAA or GDPR, *PriCell* brings unprecedented value for the healthcare domain, exemplified in this work for single-cell analysis, where the data are scarce and sparse. These benefits are extensible to federated healthcare scenarios that rely on machine learning, and constitute an important landmark for real-world applications of collaborative training between healthcare institutions while preserving privacy.

EXPERIMENTAL PROCEDURES

Resource availability

Lead contact

Further information and requests for resources should be directed to the lead contact, Sinem Sav (sinem.sav@epfl.ch).

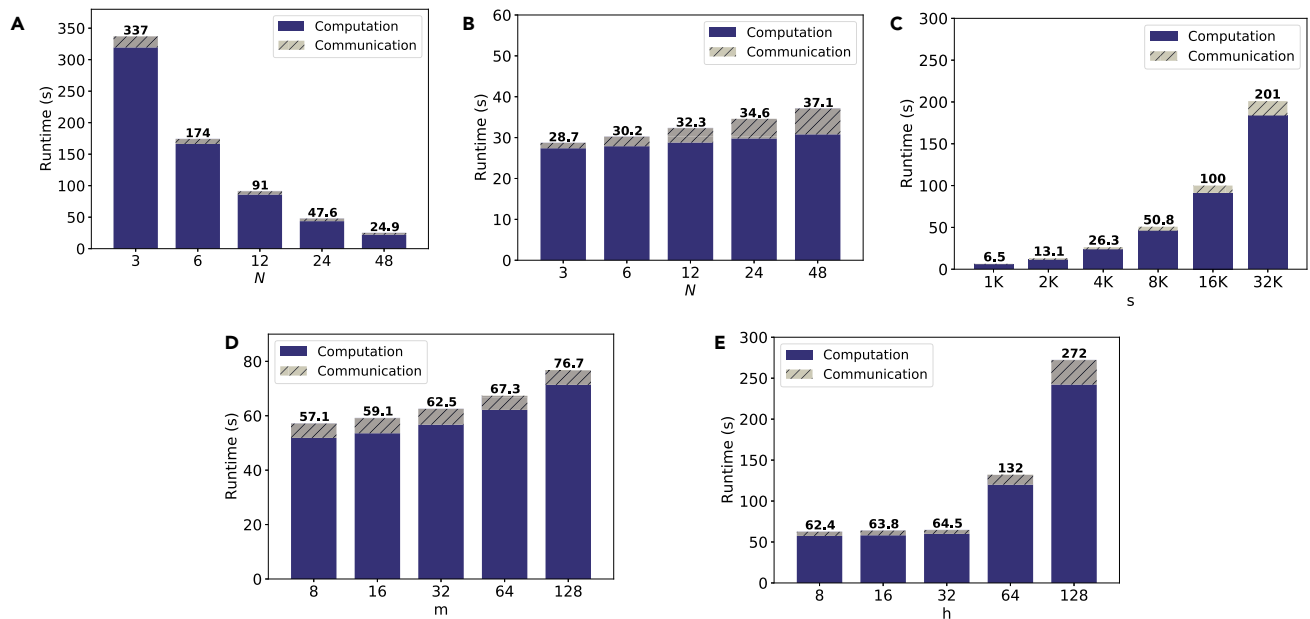


Figure 5. PriCell's training execution time and communication overhead for one training epoch with increasing number of parties, data samples, features, and filters

The computation is single-threaded in a virtual network with an average network delay of 0.17 ms and 1 Gbps bandwidth on 10 Linux servers with an Intel Xeon E5-2680 v.3 CPUs running at 2.5 GHz with 24 threads on 12 cores and 256 GB RAM.

- (A) Increasing number of parties (N) when the number of global data samples (s) is fixed to 18,000.
 (B) Increasing number of parties (N), each having 500 samples.
 (C) Increasing number of data samples (s) when $N = 10$.
 (D) Increasing number of features (m) when $N = 10$.
 (E) Increasing number of filters (h) when $N = 10$.

Materials availability

No new biological materials were generated by this study.

Data and code availability

This study uses previously published datasets.^{1–3} We refer to the reader to Arvaniti and Claassen⁴ (<https://zenodo.org/record/5597098#.YXbaz9ZBzt0>) to obtain the CMV and AML datasets, and to Galli et al.³ for the repository including the samples from healthy, NIND, and RRMS donors (<http://flowrepository.org/experiments/2166/>). Our implementation for this study have been deposited at Zenodo (<https://doi.org/10.5281/zenodo.6330988>).

System and threat model

In this section, we detail PriCell's system and threat model, which is based on the state-of-the-art privacy-preserving framework.³⁴ In PriCell's scenario, there are N healthcare institutions (parties), each holding its own patient dataset and collectively training a neural network model, without sharing/transferring their local data. Our aim is to preserve the confidentiality of the local data, the intermediate model updates in the FL setting, the querier's evaluation data, and optionally the final model (see Figure 1 for the system overview). We rely on a synchronous learning protocol, assuming that all parties are available throughout the training and evaluation executions. We note here that this assumption can be relaxed by using different HE schemes, such as threshold or multi-key HE,^{42,43} but with a relaxed security assumption for the former and an increased computation cost for the latter.

We consider an $N - 1$ passive-adversary threat model. Hence, we assume that all parties follow the protocol, and up to $N - 1$ colluding parties must not be able to extract any information about the model or other party's input data. We rely on MHE to meet these confidentiality requirements under the posed threat model. In the following section, we briefly introduce the background on the used MHE scheme.

Multiparty Homomorphic Encryption (MHE)

Here, we give the fundamentals of MHE and describe the most important cryptographic details. Note that we use italic fonts for the cryptographic terms when first introduced and these terms are explained in Note S1, cryptography glossary.

We rely on a variant of the Cheon-Kim-Kim-Song (CKKS)⁴⁴ cryptographic scheme that is based on the *ring* learning with errors (RLWE) problem (also post-quantum secure⁴⁵) and that provides approximate arithmetic over vectors of complex numbers, i.e., $\mathbb{C}^{N/2}$ (a complex vector of $N/2$ slots). Encrypted operations over $\mathbb{C}^{N/2}$ can be carried out in a SIMD fashion, which allows for excellent amortization. Therefore, adopting an efficient packing strategy that maximizes the usage of all the slots has a significant impact on the overall computation time.

Mouchet et al.⁴⁶ show how to construct a threshold variant of RLWE schemes, where the parties have their own secret key and collaborate to establish the collective public keys. In this setting, a model can be collectively trained without having to share the individual secret keys of the parties; this prevents the parties' decryption functionality, without the collaboration of all the parties.

Note that a fresh CKKS ciphertext permits only a limited number of homomorphic operations to be carried out. To enable further homomorphic operations, a ciphertext must be refreshed with a *bootstrapping operation* once it is exhausted. This operation is a costly function and requires communication in our system; hence we optimize the circuit to minimize the number of bootstrapping operations and the number of ciphertexts to be bootstrapped (see detailed neural network circuit in this section).

Datasets

We detail the features of the three used datasets:

Non-inflammatory neurological disease (NIND), relapsing–remitting multiple sclerosis (RRMS)

We rely on a large cohort of peripheral blood mononuclear cells, including 29 healthy donors (HD), 31 NIND, and 31 RRMS donors.³ The dataset

comprises samples with a varying number of cells for each donor and 35 markers for each cell. We use this dataset for two classification tasks: (1) HD versus NIND and (2) HD versus RRMS, as shown in Figures 3 and 4. For both NIND and RRMS experiments, and in all experimental settings, we use 48 donors (24 HD, 24 NIND/RRMS) for training and 12 donors (5 HD, 7 NIND/RRMS) for testing.

Cytomegalovirus Infection (CMV)

We use a mass cytometry dataset¹ for the classification of CMV. This dataset comprises samples from 20 donors with a varying number of cells for each donor and mass cytometry measurements of 37 markers for each cell, and has 11 CMV– and 9 CMV+ labels. We use 14 donors for training and 6 donors as a test set in all experimental settings.

Acute Myeloid Leukaemia (AML)

We rely on the mass cytometry dataset from Levine et al.² for the three-class classification problem for healthy, cytogenetically normal (CN), and core-binding factor translocation (CBF). For each cell, the dataset includes mass cytometry measurements of 16 markers. As in the original work,⁴ we use the AML samples with at least 10% CD34+ blast cells with the availability of additional cytogenetic information. The final training dataset comprises three healthy bone marrows (BM1, BM2, and BM3), two CN samples (SJ10 and SJ12), and two CBF samples (SJ1 and SJ2). The test set in all experimental settings comprises two healthy bone marrows (BM4, BM5), one CN (SJ13), and three CBF (SJ3, SJ4, and SJ5) samples.

The individual donors in all aforementioned training sets are then evenly distributed among N parties for *PriCell* collective training. To construct our baselines and to make a fair comparison with the baseline, we use the same data preprocessing for all experiments per setting (centralized CellCnn, Local, or *PriCell*). We give the details of the data preprocessing and parameter selection, in Note S3.

Local neural network operations

In this section, we give a high-level description of the neural network circuit that is evaluated in the encrypted domain (a detailed and step-by-step description can be found in Note S4). We list the frequently used symbols and notations in Note S2 and Table S1.

We first present the changes introduced to the original CellCnn circuit to enable an efficient evaluation under encryption: (1) we approximate the non-polynomial activation functions by polynomials by using least-squares approximation, (2) we replace the max pooling with an average pooling, and (3) we replace the ADAM optimizer with the stochastic gradient descent (SGD) optimizer with mean-squared error and momentum acceleration. Finally, we introduce the packing strategy used in *PriCell* and give a high-level circuit overview. We give more details on these steps and optimizations in Note S4, and we empirically evaluate the effect of these optimizations on the model accuracy in a distributed setting in the results section.

Polynomial approximations

With additions and multiplications, the CKKS scheme can efficiently evaluate polynomials in the encrypted domain. However, these two basic operations are not sufficient for easily evaluating non-linear functions, such as sign or sigmoid. A common strategy to circumvent this problem is to find a polynomial approximation of the desired function. We rely on polynomial least-squares approximations for the non-polynomial activation functions, such as sigmoid, and we use identity function after convolution (instead of ReLU). We show in our results section that these changes only have a negligible effect on the model accuracy.

Pooling

The original CellCnn circuit makes use of both max pooling and average pooling. Max pooling requires the computation of the non-linear sign function which cannot be efficiently done under encryption. We replace the max pooling with the average pooling, which is a linear transformation and brings the following advantages: (1) it is efficient for computing under encryption with only additions and constant multiplication, (2) it simplifies the backward pass under encryption, and (3) it commutes with other linear transformations or functions, such as the convolution and the identity activation, which allows for an efficient preprocessing of the data and reduces the online execution cost. Indeed, we are able to pre-compute the average pooling on the data, which reduces the input size of a batch of samples from $n \times c \times m$ to $n \times m$, i.e., we remove the dependency on c .

Optimizer

The original CellCnn training relies on the ADAM optimizer, which requires the computation of square roots and inverses. Although approximating these operations is possible, a high-precision approximation requires an excessive use of ciphertext levels and significantly reduces the efficiency of the training. To avoid these costly operations, we rely instead on the SGD optimizer with momentum acceleration that, for an equivalent amount of epochs, shows a comparable rate of convergence to the ADAM optimizer.

Packing strategy

The CKKS scheme provides complex arithmetic on $\mathbb{C}^{N/2}$ in a SIMD fashion. The native operations are addition, multiplication by a constant, multiplication by a plaintext, multiplication by a ciphertext, *slots rotation* (shifting the values in the vector), and complex conjugation. As the rotations are expensive, when considering encrypted matrix operations one of the main challenges is to minimize the number of rotations, which can be done by adopting efficient packing strategies and algorithms. We give more details about the packing and algorithms in Note S4.

With the aforementioned pre-computed pooling, only a $1 \times m$ vector is needed to represent a sample, instead of a $c \times m$ matrix. Hence, we pack an entire batch of n samples in a single ciphertext and compute the forward and backward pass on the whole batch in parallel, which reduces the complexity of the training proportionally to the size of the batch.

Encrypted circuit overview

Given a batch size of n samples, each sample being a matrix $L_{c \times m}$, for c the number of cells per sample and m number of features (markers) per cell, we first evaluate the mean pooling across the cells in plaintext. The result is a set of n vectors of size $1 \times m$, which is packed in an $L_{n \times m}$ matrix. The 1D convolution is evaluated with an $L_{n \times m} \times C_{m \times k}$ matrix multiplication. We feed the result to the dense layer $W_{k \times o}$, where o is the number of output labels. Finally, we perform an approximated activation function to the output of the dense layer. Our encrypted circuit with reduced complexity is

$$Y_{n \times o} = \text{Poly}_{\text{act}}((\text{MeanPool}(L_{n \times c \times m}) \times C_{m \times k}) \times W_{k \times o}).$$

Experimental settings

We implemented our solution in Go (Go Programming Language, <https://go.dev/>) by using the open-source lattice-based cryptography Lattigo (a library for lattice-based HE in Go, <https://github.com/tuneinsight/lattigo>). We use the implementation of CellCnn⁴ to preprocess the data and to construct baselines. We use Onet (Cothority Network Library, <https://github.com/dedis/onet>) to build a decentralized system and Mininet (<http://mininet.org>) to evaluate our system in a virtual network with an average network delay of 0.17 ms and 1 Gbps bandwidth on 10 Linux servers with Intel Xeon E5-2680 v.3 CPUs running at 2.5 GHz with 24 threads on 12 cores and 256 GB RAM. The parties communicate over TCP with secure channels (TLS). We choose security parameters that achieve security of at least 128 bits.⁴⁷

SUPPLEMENTAL INFORMATION

Supplemental information can be found online at <https://doi.org/10.1016/j.patter.2022.100487>.

ACKNOWLEDGMENTS

We would like to thank Apostolos Pyrgelis, David Froelicher, and Sylvain Châtel who gave valuable feedback on the manuscript. We also thank Shufan Wang and Joao Sa Sousa for their contribution on the experiments and benchmarking. This work was partially supported by grant no. 2017-201 of the Strategic Focal Area “Personalized Health and Related Technologies (PHRT)” of the ETH Domain. M.C. is a member of the Machine Learning Cluster of Excellence, EXC no. 2064/1 – project no. 390727645.

AUTHOR CONTRIBUTIONS

S.S., J.R.T.-P., M.C., and J.-P.H. conceived the study. S.S. and J.-P.B. developed the methods, implemented them, and performed the experiments. All authors contributed to the methodology and wrote the manuscript.

DECLARATION OF INTERESTS

J.R.T.-P. and J.-P.H. are co-founders of the start-up Tune Insight SA (<https://tuneinsight.com>). The other authors declare no competing interests.

The methods presented in this work are partially covered by the PCT patent applications under publication nos. WO/2022/042848⁴⁸ and WO/2021/223873.⁴⁹

Received: December 18, 2021

Revised: February 14, 2022

Accepted: March 14, 2022

Published: April 18, 2022

REFERENCES

- Horowitz, A., Strauss-Albee, D., Leipold, M., Kubo, J., Nemat-Gorgani, N., Dogan, O., Dekker, C.L., Mackey, S., Maecker, H., Swan, G.E., and Davis, M.M. (2013). Genetic and environmental determinants of human nk cell diversity revealed by mass cytometry. *Sci. Transl. Med.* 5, 208ra145. <https://doi.org/10.1126/scitranslmed.3006702>.
- Levine, J., Simonds, E., Bendall, S., Davis, K., Amir, E.A., Tadmor, M., et al. (2015). Data-driven phenotypic dissection of aml reveals progenitor-like cells that correlate with prognosis. *Cell* 162, 184–197. <https://doi.org/10.1016/j.cell.2015.05.047>.
- Galli, E., Hartmann, F.J., Schreiner, B., Ingelfinger, F., Arvaniti, E., Diebold, M., Mrdjen, D., van der Meer, F., Krieg, C., Nimer, F.A., and Sanderson, N. (2019). GM-CSF and CXCR4 define a t helper cell signature in multiple sclerosis. *Nat. Med.* 25, 1290–1300. <https://doi.org/10.1038/s41591-019-0521-4>.
- Arvaniti, E., and Claassen, M. (2017). Sensitive detection of rare disease-associated cell subsets via representation learning. *Nat. Commun.* 8, 14825. <https://doi.org/10.1038/ncomms14825>.
- Wang, T., Bai, J., and Nabavi, S. (2021). Single-cell classification using graph convolutional networks. *BMC Bioinf.* 22, 364. <https://doi.org/10.1186/s12859-021-04278-2>.
- Kirby, S., Eng, P., Danter, W., George, C., Francovic, T., Ruby, R.R., and Ferguson, K.A. (1999). Neural network prediction of obstructive sleep apnea from clinical criteria. *Chest* 116, 409–415. <https://doi.org/10.1378/chest.116.2.409>.
- Vieira, S., Pinaya, W.H., and Mechelli, A. (2017). Using deep learning to investigate the neuroimaging correlates of psychiatric and neurological disorders: methods and applications. *Neurosci. Biobehav. Rev.* 74, 58–75. <https://doi.org/10.1016/j.neubiorev.2017.01.002>.
- Uddin, M., Wang, Y., and Woodbury-Smith, M. (2019). Artificial intelligence for precision medicine in neurodevelopmental disorders. *NPJ Digital Med.* 2. <https://doi.org/10.1038/s41746-019-0191-0>.
- Kaissis, G., Makowski, M., Rückert, D., and Braren, R. (2020). Secure, privacy-preserving and federated machine learning in medical imaging. *Nat. Machine Intelligence* 2, 305–311. <https://doi.org/10.1038/s42256-020-0186-1>.
- Regev, A., Teichmann, S., Lander, E., Amit, I., Benoist, C., Birney, E., Bodenmiller, B., Campbell, P., Carninci, P., Clatworthy, M., et al. (2017). Science forum: the human cell atlas. *Elife* 6, e27041. <https://doi.org/10.7554/eLife.27041>.
- McMahan, B., Moore, E., Ramage, D., Hampson, S., and Arcas, B.A. (2017). Communication-efficient learning of deep networks from decentralized data. In Proceedings of the 20th International Conference on Artificial Intelligence and Statistics, 54, A. Singh and J. Zhu, eds. (Proceedings of Machine Learning Research (PMLR)), pp. 1273–1282. <https://proceedings.mlr.press/v54/mcmahan17a.html>.
- Sadilek, A., Liu, L., Nguyen, D., Kamruzzaman, M., Serghiou, S., Rader, B., Ingeman, A., Mellem, S., Kairouz, P., Nsoosis, E.O., et al. (2021). Privacy-first health research with federated learning. *NPJ Digital Med.* 4, 132. <https://doi.org/10.1038/s41746-021-00489-2>.
- Sheller, M., Edwards, B., Reina, G., Martin, J., Pati, S., Kotrotsou, A., Milchenko, M., Xu, W., Marcus, D., Colen, R.R., and Bakas, S. (2020). Federated learning in medicine: facilitating multi-institutional collaborations without sharing patient data. *Sci. Rep.* 10, 12598. <https://doi.org/10.1038/s41598-020-69250-1>.
- Gaye, A., Marcon, Y., Kutschke, J., Laflamme, P., Turner, A., Jones, E., Minion, J., Boyd, A.W., Newby, C.J., Nuotio, M.-L., et al. (2014). Datashield: taking the analysis to the data, not the data to the analysis. *Int. J. Epidemiol.* 43, 1929–1944. <https://doi.org/10.1093/ije/dyu188>.
- Warnat-Herresthal, S., Schultze, H., Shastry, K., Manamohan, S., Mukherjee, S., Garg, V., Sarveswara, R., Händler, K., Pickkers, P., Aziz, N.A., et al. (2021). Swarm learning for decentralized and confidential clinical machine learning. *Nature* 594, 265–270. <https://doi.org/10.1038/s41586-021-03583-3>.
- Melis, L., Song, C., De Cristofaro, E., and Shmatikov, V. (2019). Exploiting unintended feature leakage in collaborative learning. In 2019 IEEE Symposium on Security and Privacy (SP), pp. 691–706. <https://doi.org/10.1109/SP.2019.00029>.
- Nasr, M., Shokri, R., and Houmansadr, A. (2019). Comprehensive privacy analysis of deep learning: passive and active white-box inference attacks against centralized and federated learning. In 2019 IEEE Symposium on Security and Privacy (SP) (IEEE), pp. 739–753. <https://doi.org/10.1109/SP.2019.00065>.
- Hitaj, B., Ateniese, G., and Perez-Cruz, F. (2017). Deep models under the gan: information leakage from collaborative deep learning. In Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security. CCS '17 (Association for Computing Machinery), pp. 603–618. ISBN 9781450349468. <https://doi.org/10.1145/3133956.3134012>.
- Wang, Z., Mengkai, S., Zhang, Z., Song, Y., Wang, Q., and Qi, H. (2019). Beyond Inferring Class Representatives: User-Level Privacy Leakage from Federated Learning (IEEE), pp. 2512–2520. <https://doi.org/10.1109/INFOCOM.2019.8737416>.
- Zhu, L., Liu, Z., and Han, S. (2019). Deep Leakage from Gradients (Curran Associates Inc.). <http://papers.nips.cc/paper/9617-deep-leakage-from-gradients>.
- Choudhury, O., Gkoulalas-Divanis, A., Salonidis, T., Sylla, I., Park, Y., Hsu, G., and Das, A. (2019). Differential privacy-enabled federated learning for sensitive health data. <https://arxiv.org/abs/1910.02578>.
- Kim, M., Lee, J., Ohno-Machado, L., and Jiang, X. (2020). Secure and differentially private logistic regression for horizontally distributed data. *IEEE Trans. Inf. Forensics Secur.* 15, 695–710. <https://doi.org/10.1109/TIFS.2019.2925496>.
- Li, W., Milletari, F., Xu, D., Rieke, N., Hancox, J., Zhu, W., Baust, M., Cheng, Y., Ourselin, S., Cardoso, M.J., and Feng, A. (2019). Privacy-preserving federated brain tumour segmentation. In International Workshop in Machine Learning in Medical Imaging (MLMI) (Springer). https://doi.org/10.1007/978-3-030-32692-0_16.
- Jayaraman, B., and Evans, D. (2019). Evaluating differentially private machine learning in practice. In 28th USENIX Security Symposium (USENIX Security 19) (USENIX Association), pp. 1895–1912, ISBN 978-1-939133-06-9. <https://www.usenix.org/conference/usenixsecurity19/presentation/jayaraman>.
- Jagadeesh, K.A., Wu, D.J., Birgmeier, J.A., Boneh, D., and Bejerano, G. (2017). Deriving genomic diagnoses without revealing patient genomes. *Science* 357, 692–695. <https://doi.org/10.1126/science.aam9710>.
- Cho, H., Wu, D., and Berger, B. (2018). Secure genome-wide association analysis using multiparty computation. *Nat. Biotechnol.* 36, 547–551. <https://doi.org/10.1038/nbt.4108>.
- Constable, S., Tang, Y., Wang, S., Jiang, X., and Chapin, S. (2015). Privacy-preserving gwas analysis on federated genomic datasets. *BMC Med. Inf. Decis. Making* 15, S2. <https://doi.org/10.1186/1472-6947-15-S5-S2>.
- Kamm, L., Bogdanov, D., Laur, S., and Vilo, J. (2013). A new way to protect privacy in large-scale genome-wide association studies. *Bioinformatics* 29, 886–893. <https://doi.org/10.1093/bioinformatics/btt066>.

29. Hie, B., Cho, H., and Berger, B. (2018). Realizing private and practical pharmacological collaboration. *Science* 362, 347–350. <https://doi.org/10.1126/science.aat4807>.
30. Kim, M., Song, Y., Wang, S., Yuhou, X., and Jiang, X. (2018a). Secure logistic regression based on homomorphic encryption: design and evaluation. *JMIR Med. Inform.* 6, e19. <https://doi.org/10.2196/medinform.8805>.
31. Bonte, C., and Vercauteren, F. (2018). Privacy-preserving logistic regression training. *BMC Med. Genomics* 11, 86. <https://doi.org/10.1186/s12920-018-0398-y>.
32. Froelicher, D., Troncoso-Pastoriza, J.R., Raisaro, J.L., Cuendet, M.A., Sousa, J.S., Cho, H., Berger, B., Fellay, J., and Hubaux, J.-P. (2021a). Truly privacy-preserving federated analytics for precision medicine with multiparty homomorphic encryption. *Nat. Commun.* 12, 5910. <https://doi.org/10.1038/s41467-021-25972-y>.
33. Raisaro, J.L., Troncoso-Pastoriza, J.R., Misbach, M., Sousa, J.S., Pradervand, S., Missiaglia, E., Michielin, O., Ford, B., and Hubaux, J.P. (2019). Medco: enabling secure and privacy-preserving exploration of distributed clinical and genomic data. *IEEE/ACM Trans. Comput. Biol. Bioinformatics* 16, 1328–1341. <https://doi.org/10.1109/TCBB.2018.2854776>.
34. Sav, S., Pyrgelis, A., Troncoso-Pastoriza, J.R., Froelicher, D., Bossuat, J.P., Sousa, J.S., and Hubaux, J.P. (2021). Poseidon: privacy-preserving federated neural network learning. In *Network and Distributed System Security Symposium (NDSS)*. <https://doi.org/10.14722/ndss.2021.24119>.
35. Rawat, W., and Wang, Z. (2017). Deep convolutional neural networks for image classification: a comprehensive review. *Neural Comput.* 29, 1–98. https://doi.org/10.1162/NECO_a_00990.
36. Rahman, M.A., Rahman, T., Laganière, R., and Mohammed, N. (2018). Membership inference attack against differentially private deep learning model. *Trans. Data Privacy* 11, 61–79. <https://www.tdp.cat/issues16/tdp.a289a17.pdf>.
37. Kim, J.W., Jang, B., and Yoo, H. (2018b). Privacy-preserving aggregation of personal health data streams. *PLoS One* 13, 1–15. <https://doi.org/10.1371/journal.pone.0207639>.
38. Chen, W., Sotiraki, K., Chang, I., Kantarcioglu, M., and Popa, R.A. (2021). HOLMES: a platform for detecting malicious inputs in secure collaborative computation. <https://eprint.iacr.org/2021/1517>.
39. Yang, R., Au, M.H., Zhang, Z., Xu, Q., Yu, Z., and Whyte, W. (2019). Efficient Lattice-Based Zero-Knowledge Arguments with Standard Soundness: Construction and Applications, pp. 147–175. ISBN 978-3-030-26947-0. https://doi.org/10.1007/978-3-030-26948-7_6.
40. Baum, C., and Nof, A. (2020). Concretely-efficient Zero-Knowledge Arguments for Arithmetic Circuits and Their Application to Lattice-Based Cryptography. In *Public-Key Cryptography – PKC 2020: 23rd IACR International Conference on Practice and Theory of Public-Key Cryptography*, Edinburgh, UK, May 4–7, 2020, *Proceedings, Part I* (Springer-Verlag), pp. 495–526, ISBN 978-3-030-45373-2.
41. Kairouz, P., McMahan, H.B., Avent, B., Bellet, A., Bennis, M., Bhagoji, A.N., Bonawitz, K., Charles, Z., Cormode, G., Cummings, R., and D'Oliveira, R.G. (2021). Advances and open problems in federated learning. *Found. Trends Machine Learn.* 14, 1–210. <https://doi.org/10.1561/22000000083>.
42. López-Alt, A., Tromer, E., and Vaikuntanathan, V. (2012). On-the-fly multiparty computation on the cloud via multikey fully homomorphic encryption. In *Proceedings of the Forty-Fourth Annual ACM Symposium on Theory of Computing (STOC '12; Association for Computing Machinery)*, pp. 1219–1234. ISBN 9781450312455. <https://doi.org/10.1145/2213977.2214086>.
43. Shamir, A. (1979). How to share a secret. *Commun. ACM* 22, 612–613. <https://doi.org/10.1145/359168.359176>.
44. Cheon, J.H., Kim, A., Kim, M., and Song, Y. (2017). Homomorphic encryption for arithmetic of approximate numbers. In *Springer International Conference on the Theory and Application of Cryptology and Information Security (ASIACRYPT)*. https://doi.org/10.1007/978-3-319-70694-8_15.
45. Acar, A., Aksu, H., Uluagac, A.S., and Conti, M. (2018). A survey on homomorphic encryption schemes: theory and implementation. *ACM Comput. Surv.* 51, 1–35. <https://doi.org/10.1145/3214303>.
46. Mouchet, C., Troncoso-pastoriza, J.R., Bossuat, J.P., and Hubaux, J.P. (2021). Multiparty Homomorphic Encryption from Ring-Learning-With-Errors (PETS). <https://doi.org/10.2478/popets-2021-0071>.
47. Albrecht, M., Chase, M., Chen, H., Ding, J., Goldwasser, S., Gorbunov, S., Halevi, S., Hoffstein, J., Laine, K., Lauter, K., and Lokam, S. (2018). Homomorphic encryption security standard. *Tech. Rep. (Homomorphic Encryption.org)*.
48. Sav, S., Troncoso-Pastoriza, J.R., Pyrgelis, A., Froelicher, D., Gomes de Sá e Sousa, J.A., Bossuat, J.P., and Hubaux, J.-P. (2022). System and Method for Privacy-Preserving Distributed Training of Neural Network Models on Distributed Datasets (EPFL).
49. Froelicher, D., Troncoso-Pastoriza, J.R., Pyrgelis, A., Sav, S., Gomes de Sá e Sousa, J.A., Hubaux, J.P., and Bossuat, J.P. (2021). System and Method for Privacy-Preserving Distributed Training of Machine Learning Models on Distributed Datasets (EPFL).