



HAL
open science

A cautionary tale on the cost-effectiveness of collaborative AI in real-world medical applications

Francesco Cremonesi, Lucia Innocenti, Sebastien Ourselin, Vicky Goh, Michela
Antonelli, Marco Lorenzi

► **To cite this version:**

Francesco Cremonesi, Lucia Innocenti, Sebastien Ourselin, Vicky Goh, Michela Antonelli, et al.. A cautionary tale on the cost-effectiveness of collaborative AI in real-world medical applications. 2025. <hal-04893012>

HAL Id: hal-04893012

<https://hal.science/hal-04893012v1>

Preprint submitted on 17 Jan 2025

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



HAL Authorization

A cautionary tale on the cost-effectiveness of collaborative AI in real-world medical applications.

Francesco Cremonesi^{1†}, Lucia Innocenti^{1,2†}, Sebastien Ourselin²,
Vicky Goh², Michela Antonelli^{2‡}, Marco Lorenzi^{1*‡}

¹Epione Research Group, Inria Center of Université Côte d’Azur, France.

²King’s College London, London, UK.

*Corresponding author(s). E-mail(s): marco.lorenzi@inria.fr;

†joint first author

‡joint last author

Abstract

Background. Federated learning (FL) has gained wide popularity as a collaborative learning paradigm enabling collaborative AI in sensitive healthcare applications. Nevertheless, the practical implementation of FL presents technical and organizational challenges, as it generally requires complex communication infrastructures. In this context, consensus-based learning (CBL) may represent a promising collaborative learning alternative, thanks to the ability of combining local knowledge into a federated decision system, while potentially reducing deployment overhead.

Methods. In this work we propose an extensive benchmark of the accuracy and cost-effectiveness of a panel of FL and CBL methods in a wide range of collaborative medical data analysis scenarios. The benchmark includes 7 different medical datasets, encompassing 3 machine learning tasks, 8 different data modalities, and multi-centric settings involving 3 to 23 clients.

Findings. Our results reveal that CBL is a cost-effective alternative to FL. When compared across the panel of medical dataset in the considered benchmark, CBL methods provide equivalent accuracy to the one achieved by FL. Nonetheless, CBL significantly reduces training time and communication cost (resp. 15 fold and 60 fold decrease) ($p < 0.05$).

Interpretation. This study opens a novel perspective on the deployment of collaborative AI in real-world applications, whereas the adoption of cost-effective methods is instrumental to achieve sustainability and democratisation of AI by alleviating the need for extensive computational resources.

Funding. This work was supported by the 3IA Côte d’Azur Investments in

the Future project managed by the National Research Agency (ANR-15-IDEX-01 and ANR-19-P3IA-0002), by the grants TRAIN-ANR-22-FAI1-0003, and Fed-BioMed-ANR-19-CE45-0006.

Keywords: Collaborative learning, healthcare, sustainable AI, trustworthy AI, federated learning consensus-based learning, medical imaging,

1 Introduction

Collaborative learning (CL) has become a popular paradigm for the exploitation of Artificial Intelligence (AI) in healthcare applications^{1,2}. CL allows multiple parties to jointly solve analytical tasks by combining local knowledge extracted from the respective data. These abstractions can, for instance, take the form of global data statistics, machine learning (ML) model parameters, or model’s predictions. This paradigm is rapidly gaining popularity in healthcare applications due to its appealing promises for accuracy/privacy trade-off^{3,4}.

Federated Learning (FL) has been identified as a key CL paradigm², focusing on collaboratively optimizing model parameters across clients, each holding local datasets (Figure 1, left). FL is based on an iterative optimization paradigm in which each client shares model parameters partially trained on the respective local data, which are then aggregated by a central server to obtain a global model. Although FL has demonstrated significant advantages in security and privacy over centralized approaches, its implementation in real-world applications is not straightforward and can face technical and organizational challenges⁵, reflecting the technology’s complexity as well as significant computational costs⁶.

Consensus-based learning (CBL)⁷⁻⁹ represents a valid CL alternative to FL. CBL does not rely on a shared training routine nor on a common model architecture across parties, but instead combines the predictions obtained from the different models independently trained by each client on their local data. CBL thus relies on an off-line setting, in which information is exchanged only at inference time (Figure 1, right).

Bearing in mind the duality between these paradigms, when it comes to implementing collaborative frameworks we currently lack quantitative benchmarks illustrating the quality and cost-effectiveness of specific CL settings in real-world medical applications. On the one hand, FL has been widely investigated for CL applications, in particular for healthcare¹⁰⁻¹². On the other hand, while there exists a large body of literature demonstrating the effectiveness of CBL over training local models [CITAZIONE], its applications to distributed datasets remain underinvestigated except, to our knowledge, for the work of Guha et al.⁷ (preprint) and Chaudhari et al.¹³. In the context of medical imaging applications, the capabilities of different CL paradigms have been discussed in Gupta et al.¹⁴, without, however, providing any experimental evaluation. To our knowledge, a general reference point that allows comparison between CBL and FL paradigms while considering the accuracy and cost-effectiveness of these paradigms is still missing.




Fig. 1: Training and inference phases for federated learning (FL, on the left) and consensus-based learning (CBL, on the right). In FL training is performed collaboratively to produce a common global model across clients. The global model is subsequently used for inference on new data instances. CBL instead requires clients to train a model on the respective local data independently. Inference on new data instances is performed collaboratively through consensus.

In this work we carry out the first comparative analysis of the capabilities and cost-effectiveness of CBL and FL. We compare the performance of a panel of state-of-the-art FL and CBL methods on a variety of learning tasks and data modalities representative of real-world medical use cases. Our results show that in most of the evaluated benchmarks, CBL is a cost-effective alternative to FL, achieving comparable performance with the advantage of reducing training and communication costs. This opens a novel perspective on the deployment of collaborative AI in the real-world, in which the opportune choice of CL paradigms suited to the specifics of each task can mitigate the implementation burden of CL.

2 Methods

This section introduces the dataset and benchmarks adopted in our work (Section 2.1 and 2.3), along with the panel of FL and CBL methods evaluated in the experiments (Section 2.2).

2.1 Benchmark design

We formalized our goal of comparing CBL and FL methods by formulating the two following research questions: i) are there particular combinations of data and ML task where one of the two paradigms provides significantly better model performance? ii) what is the difference in resources requirements between FL and CBL to achieve the same accuracy level on a given combination of data and ML task?

Table 1: Collaborative learning (CL) methods evaluated in the benchmark: six methods for federated learning (FL) and five methods for consensus-based learning (CBL).

Method	CL Paradigm	Reference
FEDAVG	FL	McMahan et. Al., 2017 ¹⁷
FEDPROX	FL	Li et Al., 2018 ¹⁸
SCAFFOLD	FL	Karimireddy et Al., 2020 ¹⁹
FEDADAM	FL	Reddi et Al., 2020 ²⁰
FEDADAGRAD	FL	Reddi et Al., 2020 ²⁰
FEDYOGI	FL	Reddi et Al., 2020 ²⁰
AVG	CBL	Guha et Al., 2019 ⁷
MV	CBL	Safdar et Al., 2021 ²¹
STAPLE	CBL	Warfield et Al., 2004 ²²
UBE	CBL	<i>Inspired by</i> Ruta et Gabrys, 2000 ²³
ABE	CBL	<i>Inspired by</i> Ruta et Gabrys, 2000 ²³

To answer these questions we identified seven heterogeneous benchmarks, defined by their datasets and associated ML tasks, considered representative of real-world case studies. The targeted tasks include segmentation of MRI (*FedProstate*, *FedIXI*, *FeTS*) and segmentation of CT images (*FedKITS*), diagnosis (*FedHeart*), disease subtyping (*FedISIC*), and survival probability prediction (*FedTCGA-BRCA*). The benchmarks present a high degree of heterogeneity in terms of the number of clients in the federation (from 3 to 23), sample sizes (100 to over 2000 patients), data modalities (total of 8 different modalities), and distribution of data among clients, thus reflecting the heterogeneity of real-world applications.

Concerning FL methods, our benchmark covers standard aggregation strategies adopted in FL benchmarks for healthcare in the literature^{10,15}. We included the standard aggregation mechanism proposed in the seminal work of McMahan et al., FEDAVG, along with subsequent approaches aimed at mitigating the impact of client heterogeneity in the optimization: FEDPROX, SCAFFOLD, as well as FEDADAM, FEDYOGI, and FEDADAGRAD.

For CBL methods, we tested several classical fusion strategies to combine the predictions of local models, such as plain averaging (AVG) and majority voting (MV). Moreover, for segmentation tasks, we used STAPLE, which optimizes consensus through expectation-maximization. Finally, we proposed two CBL methods based on decision averaging: uncertainty-based (UBE) and autoencoder-based ensembling (ABE). Both methods rely on the weighted average of local predictions. For UBE the weights are estimated by quantifying the uncertainty of the predicted labels. For ABE the uncertainty was quantified by the reconstruction error on the testing point of an autoencoder trained on the local dataset. This measure is a proxy for out-of-distribution detection¹⁶.

Table 1 presents an overview of the 6 FL and 5 CBL methods used in the benchmark. Further details on all the methods are available in Section 2.2.

We analyzed the CL methods by assessing the respective performance and cost-effectiveness. As a measure of task performance we quantified the dice score, the

balanced accuracy, and the C-Index for respectively segmentation, classification, and survival tasks. We assessed the cost-effectiveness by measuring the *training time* and *network usage* required by each CL method. The training time was estimated as the total wallclock time necessary to obtain the collaborative model. The network usage was the total data transferred on the network during training. We considered a scenario in which all clients participate in every round of FL training, and we used the same model architecture across clients for each experiment. The performance of different CL methods was evaluated on the same cross-validation partitions.

2.2 Collaborative Learning Methods

We consider a collaborative setting with M clients. A dataset \mathcal{D} belonging to client i is composed of data samples $\mathcal{D}_i = \{z_{k,i}\}_{k=1}^{N_i}$, being N_i the dataset size. We consider a model f with parameters θ , a loss function \mathcal{L} , and we denote the prediction of a data instance z by $h = f(z, \theta)$.

FL is a collaborative paradigm associated with the optimization of a loss distributed among M clients defined as $\mathcal{L}(\theta) := \sum_{i=1}^M p_i \mathcal{L}_i(\theta_i)$, where the losses of local models (θ_i) are averaged by using the weights p_i . Solving this optimization problem leads to global model parameters θ_g . We consider here a comprehensive panel of state-of-the-art optimization approaches, which are at the core of the FL literature:

- **FedAvg**¹⁷ is the backbone of FL optimization, and is based on an iterative process where, at each optimization round r , clients execute a fixed number of local stochastic gradient descent steps and send the partially optimized model θ_i^r to the server. The server averages the received models according to the weights p_i to obtain a global one, θ_g^{r+1} . The global model is then sent to the clients to initialize the next optimization round.
- **FedProx**¹⁸ tackles the problem of federated optimization with data heterogeneity across clients. This approach extends FEDAVG by introducing a proximal term to the local objective function to penalize model drift from the global optimization during local training. The proximal term is controlled by a trade-off hyperparameter.
- **SCAFFOLD**¹⁹ addresses the limitations of FEDAVG in scenarios with heterogeneous data by utilizing control variates, which effectively reduces variance and corrects for client-drift in the local updates. To achieve this, SCAFFOLD maintains a state for each client (client control variate) and the server (server control variate).
- **FedAdam**, **FedYogi**, **FedAdagrad**²⁰ are adaptations of Adam, Yogi, and Adagrad optimizers, designed to suit the federated optimization setup.

CBL, on the other hand, is a class of machine learning algorithms relying on the concept of *ensembling*, a widely-explored approach that consists of obtaining robust predictions by aggregating decisions obtained by independently trained weak predictors.

The CBL paradigm is based on a *training* phase in which M models $f(\cdot, \theta_i)$ are independently trained on separated data collections \mathcal{D}_i , by minimizing local objective functions \mathcal{L}_i . These models are subsequently collected and, for a given test data z'

at *inference* time, the predictions $h_i(z') = f(z', \theta_i)$ from all the clients’ models are computed and aggregated by applying an ensembling (or fusion) strategy:

$$h_{z'} = \text{ensembling}(\{h_i(z')\}_{i=1}^M). \quad (1)$$

Typical ensembling methods proposed in the literature are:

Majority voting (MV)²¹ is often used in classification tasks, aggregates predictions by selecting the most commonly predicted class among the experts.

Staple²² is an algorithm based on expectation-maximization which, iteratively, first computes a weighted average of each local prediction, and then assigns a performance level to each client’s segmentation, which will be used as weights for the next step.

Decision averaging (DA)²⁴ is an approach based on probabilistic principles: given different datasets $\mathcal{D}_1, \dots, \mathcal{D}_M$ and associated local models $f(\cdot, \theta_1), \dots, f(\cdot, \theta_M)$, ensembling is obtained by weighing each local model’s prediction by the probability of observing the input datapoint in the local model’s data distribution $p(z' \in \mathcal{D}_i)$. Different CBL algorithms can be obtained based on the estimation of the probability $p(z' \in \mathcal{D}_i)$. In this work, among the possible DA algorithms, we consider the following:

- **Average** (AVG) approximates $p(z' \in \mathcal{D}_i)$ as a uniform distribution. This strategy can be adopted in all tasks, and for classification and segmentation problems we consider the distribution probability obtained through a *softmax* function.
- **Uncertainty based ensembling** (UBE) approximates the probability $p(z' \in \mathcal{D}_i)$ as the uncertainty of the local model on the prediction task. UBE defines averaging weights based on the model uncertainty quantified by the total element-wise variance at inference time.
- **Autoencoder based ensembling** (ABE) computes a proxy for the probability $p(z' \in \mathcal{D}_i)$ by modeling the variability of the local dataset through autoencoders trained by each client on the respective data. At inference time, weights are defined as the reconstruction error e_i on the testing data point z .

2.3 Datasets and tasks

Table 2 summarizes the datasets used for the benchmark, while further details on the composition of the datasets, the preprocessing steps, and the training task are provided in Appendix 5.

FedProstate

For the task of prostate segmentation, we assembled a dataset composed of roughly 300 T2-magnetic resonance imaging (MRI) and segmentation masks of the whole prostate, obtained from a collection of publicly available datasets and a from a retrospective cohort acquired at the Guy St. Thomas Hospital of King’s College London. Data among federated clients was split according to the acquisition protocol (MRI with or without endorectal coil), and the scanner manufacturer.

Table 2: Datasets used for the benchmark. For each dataset, we report the number, the size of clients’ local datasets, the data modality, and the task, along with the associated metric.

Dataset	#Clients	Dataset size per client	Data Modality	Task [Metric]
FedProstate	6	[32, 23, 27, 184, 5, 36]	T2 MRI	Segmentation [DiceScore]
FedHeart	4	[303, 261, 46, 130]	Tabular Data	Classification [Accuracy]
FedIXI	3	[311, 181, 74]	T1 MRI	Segmentation [DiceScore]
FedISIC	6	[12k, 3.9k, 3.3k, 225, 819, 439]	Dermoscopy images	Classification [Accuracy]
FedTCGA-BRCA	6	[311, 196, 206, 162, 162, 51]	Tabular Data	Survival [C-Index]
FedKiTS	6	[12, 14, 12, 12, 16, 30]	CT Scans	Segmentation [DiceScore]
FeTS	23	4 to 511 <i>Details in Suppl. 5</i>	T1, T1CE, T2 & FLAIR MRIs	Segmentation [DiceScore]

FedHeart

For the task of diagnosis (binary classification for the presence or absence of heart disease) we leveraged a publicly available benchmark dataset¹⁵ comprising patient demographics, medical history, and physiological measurements, summarized in 13 tabular features from roughly 700 patients. Data was split according to the referenced publication.

FedIXI

For the task of brain segmentation we used a publicly available benchmark dataset¹⁵ composed of T1-MRI images from roughly 500 patients. Data was split according to the partitions considered in the original publication.

FedISIC

For the task of skin disease subtyping we considered a publicly available benchmark dataset²⁵ composed of roughly 20000 dermoscopy images. Data was split according to the imaging acquisition system.

FedTCGA-BRCA

For the task of survival probability prediction, we used a publicly available benchmark dataset¹⁵ including 39 clinical and genomic tabular features extracted from roughly 1000 breast cancer patients. Data was split according to the acquisition site.

FedKITS

For the task of segmenting CT images, we used a publicly available benchmark dataset¹⁵ containing CT scans and segmentation masks for both the kidney and the tumor for roughly 100 patients. Data was split according to acquisition site.

FeTS

For the brain tumor segmentation task, we considered a publicly available dataset²⁶ consisting of multiple-modality MRI (T1, T2, T1CE, FLAIR) for roughly 2000 brain tumor patients. Data was split according to the acquisition site.

2.4 Experimental Setting

We split each dataset in the benchmark into training and testing partitions. For FeTS and FedProstate, we used respectively a 4-fold and a 5-fold cross-validation while for the other datasets, following du Terrail et al.¹⁵, we ran all the experiments three times with different seeds. The obtained results are the average among the runs.

For FL, a federated infrastructure was simulated using the software FedBioMed version 4.1¹. For each FL method, the final global model was collected along with the local models independently trained by each client. The local models were used to estimate the local model performance and for generating the predictions subsequently aggregated with the CBL methods. As upper-bound for the comparison, a centralized model was trained by pooling together all the local training sets.

To ensure fairness of the comparison, for each dataset we calibrated the total number of training steps to be executed in each experiment. Specifically, a number E of epochs was defined a priori, corresponding to the number of training epochs executed by the centralized model. For the local training, each client executes $\frac{E}{M}$ epochs, being M the number of clients in that configuration. For the federated strategies, we fine-tuned the number of local SGD steps s executed at each round, while the number of rounds was defined as follows: $R = E \cdot N_T / M / B / s$, where B is the batch size and N_T is the total number of samples in the training set.

3 Results


3.1 Benchmark results

3.1.1 Performance evaluation

Figure 2 shows the distribution of testing performances across datasets.


Considering the CL methods altogether, we found that collaborative approaches yielded a better average performance than the locally trained models in five out of seven benchmarks, with the best and worst performance ratios achieved respectively for ISIC ($1.42\times$ performance increase) and Heart ($0.88\times$ performance decrease). Furthermore, CL methods had on average a poorer performance compared to the pooling of all the data in a centralized training scenario (average performance decrease of $0.83\times$). This result is in line with those reported in the literature on the effectiveness of

¹<https://fedbiomed.org/>




(a) FedProstate

(b) FedHeart




(c) FedIsic

(d) FedTGCA-BRCA



(e) FedIXI

(f) FedKiTS



(g) FETS

Fig. 2: Results obtained by centralized learning (green), local learning (blue), federated learning (orange), and consensus-based learning (brown) methods. The boxplot represents the accuracy among test sets for centralized learning, local models, and CL methods. For FeTS, local accuracy results are aggregated due to the large number of clients.

Table 3: Summary results for multiple Mann-Whitney U-tests used to assess differences between the performance of consensus-based learning (CBL) and federated learning (FL) for different types of ML task. CBL outperforms FL across all tasks, however no difference was statistically significant.

Task	Average performance ratio $\frac{\text{CBL}}{\text{FL}}$	U statistic	p-value
CLASSIFICATION	1 · 11	18	0 · 26
SEGMENTATION	1 · 02	17	0 · 79
SURVIVAL	1 · 06	16	0 · 10

both FL and CBL methods for improving the performance of individual learners while maintaining a comparable performance compared to pooled datasets [CITAZIONE].

On average, CBL approaches outperformed FL in four benchmarks (average factor $1 \cdot 18\times$), had comparable performance in the FeTS (factor $1 \cdot 03\times$) and ISIC (factor $0 \cdot 97\times$) benchmarks, and had poorer performance in the KITS benchmark (factor $0 \cdot 70\times$). However, tests for the difference in the mean performance across all benchmarks between CBL and FL methods did not achieve statistical significance (Wilcoxon signed-rank test $p = 0 \cdot 2$).

We then checked whether there could be differences at the level of task types and of individual benchmarks. CBL outperformed FL in all three task categories (classification, segmentation, survival), however none of these differences were statistically significant (see Table 3). Averaging the performance over all client test sets for a given learning method, we used seven parallel statistical tests (one per benchmark) to test whether CBL methods performed worse than FL on average (see Table 4). The only two statistically significant outcomes were in the IXI and KITS benchmarks: CBL had better performance for IXI (factor $1 \cdot 18\times$, Mann-Whitney $p = 0 \cdot 009$), but poorer performance for KITS (factor $0 \cdot 70\times$, Mann-Whitney $p = 0 \cdot 017$). It is important to note that neither result maintains statistical significance after Bonferroni correction. The ensemble of these results show that overall no single collaborative learning paradigm consistently outperforms the others across all benchmarks. For example, as already pointed out in du Terrail et al.¹⁵, FL performs worse than local learning in the FedIXI benchmark, whereas CBL achieves higher accuracy. Conversely, for FedKiTS, FL methods had better performance compared to the local models, while CBL methods had poorer one.

Furthermore, to reflect the scenario where large amounts of resources can be dedicated to identifying the best approach tailored to one’s dataset, we identified the best-performing CBL and FL method separately for each benchmark and compared their performances. The best CBL method had a better performance than the best FL method in the FedHeart (factor $1 \cdot 09\times$) and FedIXI (factor $1 \cdot 06\times$) benchmarks, while the best CBL and FL methods had comparable performances in four benchmarks, and the best CBL method performed worse than the best FL method only in the FedKITS benchmark (factor $0 \cdot 77\times$). When instead comparing the basic methods for each CL category, i.e. FedAVG for FL and AVG for CBL, we found that CBL had better performance than FedAVG in three out of seven benchmarks (average factor

Table 4: Summary results for multiple Mann-Whitney U-tests used to assess differences between the performance of consensus-based learning (CBL) and federated learning (FL) across benchmarks. The only significant differences are observed in FedIXI, where CBL is associated with better performances than FL, and FedKITS, where CBL is associated with poorer performances than FL. The statistical significance does not survive after Bonferroni correction for multiple comparisons across tasks.

Experiment	Average performance ratio ($\frac{\text{CBL}}{\text{FL}}$)	U statistic	p-value
FedProstate	1 · 12×	22 · 5	0 · 20
FedHeart	1 · 36×	19	0 · 16
FedIxi	1 · 18×	29	0 · 01
FedIsic	0 · 97×	5	0 · 17
FedTGCA-BRCA	1 · 06×	10	0 · 10
FedKiTS	0 · 70×	2	0 · 02
FeTS	1 · 03×	19	0 · 54

1 · 37×), had comparable performance in the FeTS (factor 1 · 01×) and ISIC (factor 1 · 00×) benchmarks, and had worse performance in the Prostate (factor 0 · 94×) and KiTS (factor 0 · 55×) benchmarks (see Table 5).

The average performance across all client test sets, for both commonly-used and best-performing methods, is summarized in Table 5.

Having assessed that no paradigm outperforms the other in terms of task-specific performance metrics, we now look at their cost-effectiveness. The numerical evaluation of the cost-effectiveness of CBL and FL, presented in Table 6, shows that the training of CBL methods is 8 – 30 times faster than the one required by the federated methods (avg. 15× decrease, $p < 0 · 05$, Wilcoxon Signed-Rank Test for log-differences). For the same reason, the usage of bandwidth of CBL is from 35 to 120 times lower than for FL (avg. 60× decrease, $p < 0 · 05$, Wilcoxon Signed-Rank Test for log-differences).

4 Discussion

Our results show that for typical medical data analysis tasks, CBL leads to model performance on par with FL, while requiring a fraction of computational costs and bandwidth usage. This result entails profound implications for the real-world adoption of AI in healthcare.

The key advantage of CBL relies on its asynchronous training approach, thus simplifying the collaboration among partners as hospitals simply need to train their models locally to join the collaboration. Moreover, the withdrawal (or introduction) of a participant can be simply achieved by removing (or adding) their local model without the need for complex procedures or re-training. The complexity of local models can also be adapted to the availability of local resources, thus mitigating the problem of hardware heterogeneity in federated setups, and promoting the adoption of AI in healthcare. In FL, besides tuning the hyperparameters of the local models, it is also often necessary to tune method-specific hyperparameters, which increases the task

Table 5: Summary performance metrics for commonly-used (FedAVG and AVG) and best-performing training methods within the CBL and FL groups, computed separately for each benchmark. Neither of the two paradigms yields better performance in a consistent and statistically significant way.

Experiment	Paradigm	Best method	Metric	Common method	Metric
FeTS - Stats	CBL	MV	0.67	AVG	0.66
	FL	Scaffold	0.67	FedAVG	0.66
Heart	CBL	UBE	0.78	AVG	0.64
	FL	FedAdam	0.72	FedAVG	0.36
ISIC	CBL	AVG	0.62	AVG	0.62
	FL	FedYogi	0.65	FedAVG	0.62
IXI	CBL	MV	0.96	AVG	0.90
	FL	FedAdagrad	0.90	FedAVG	0.75
KITS	CBL	MV	0.42	AVG	0.29
	FL	FedProx	0.54	FedAVG	0.52
Prostate	CBL	UBE	0.81	AVG	0.75
	FL	Scaffold	0.82	FedAVG	0.80
TGCA-BRCA	CBL	AVG	0.79	AVG	0.79
	FL	FedAdagrad	0.78	FedAVG	0.70

Table 6: Comparison of training time and bandwidth usage between federated learning (FL) and consensus-based learning (CBL). For both metrics, CBL is far less costly than FL.

	Training Time [min]			Bandwidth [MB]		
	FL	CBL	FL Increment	FL	CBL	FL Increment
FedProstate	1.2×10^3	1.1×10^2	$\times 8$	4.0×10^3	1.4×10^2	$\times 35$
FedHeart	0.1	2.0×10^{-2}	$\times 9$	3.2×10^1	9.0×10^{-2}	$\times 38$
FedIXI	5.5	6.4×10^{-1}	$\times 8$	1.5×10^2	3.0	$\times 50$
FedIsic	1.9×10^2	1.0×10^1	$\times 18$	6.8×10^3	1.9×10^2	$\times 58$
FedTCGA	4.5×10^{-1}	2×10^{-2}	$\times 26$	0.9	5.42	$\times 36$
FedKITS	5.4×10^1	$6.4 \times 1.5 \times 10^2$	$\times 30$	8.8×10^4	7.4×10^2	$\times 120$
FeTS	3.4×10^4	2.9×10^3	$\times 12$	3.7×10^4	4.5×10^2	$\times 83$

complexity compared to CBL. Using the same hyperparameters on all FL methods can lead to variability in the accuracy across methods. In contrast, CBL methods exhibit lower variance among results.

Data heterogeneity, a common issue in healthcare data, may also play a role in the choice of CBL over FL approaches. Indeed, while the generalization properties of CBL may improve when ensembling models trained on heterogeneous datasets, it is known that the distributed optimization procedure of FL suffers from data heterogeneity, which causes degradation in convergence and accuracy. Such difference suggests that

adopting CBL could be beneficial in collaborations among hospitals with heterogeneous data, reflecting for example varying geographical location or image acquisition techniques.

The analysis addressed in this paper focuses on the cost-effectiveness of CBL to boost the deployment of AI in real-life healthcare applications. However, CBL does not necessarily provide quantifiable privacy guarantees. While cryptographic approaches such as secure aggregation, multi-party computation, and differential privacy are widely investigated in FL, the analysis of CBL from a privacy-preserving perspective is less explored, with some studies showing that privacy-preserving techniques can be adopted for CBL, with a potentially smaller impact on the final model accuracy when compared to FL¹³.

With this work, we hope to raise awareness of the importance of cost-effective collaborative learning paradigms in the real-world deployment of AI models on medical tasks. Starting from the empirical results presented here, additional work should be devoted to developing quantitative measures of a collaborative system’s effectiveness before deploying its infrastructure. Such a theory would allow the identification of optimal collaborative paradigms based on accounting for several aspects, such as client availability, data heterogeneity, and security requirements. Moreover, in evaluating the cost-effectiveness, estimation of the CO2 emissions and energy cost of the training would bring added value to the analysis.

Acknowledgements. LI and ML are supported by the 3IA Côte d’Azur Investments in the Future project managed by the National Research Agency (ref.n ANR-15-IDEX-01 and ANR-19-P3IA-0002). ML is funded by the grants (TRAIN - ANR-22-FAI1-0003) and (Fed-BioMed - ANR-19-CE45-0006).

Supplementary Information


5 Details on Federated Datasets, Architectures, and Hyperparameters

5.1 FedProstate

The FedProstate dataset is the federated version of the 3 major publicly available datasets on prostate cancer imaging analysis, and of 1 private dataset:

- **Medical Segmentation Decathlon - Prostate**²⁷ provides 32 prostate MRIs for training.
- **Promise12**²⁸ consists of 50 training cases obtained with different scanners. Of those, 27 cases were acquired by using an endorectal coil.
- **ProstateX**²⁹ contains prostate MRIs acquired by using two different scanners (Skyra and Triotim, both from Siemens). Segmentations of 194 cases are available³⁰.
- **Guy St. Thomas Hospital dataset (King’s College London)**. This dataset is composed by 36 MRIs acquired during the clinical routine for patients with prostate cancer under active surveillance treatment. Images were acquired with a Siemens

Fig. 3: Examples of images from the FedProstate dataset showing the heterogeneity among different clients.



Aera scanner and an expert radiologist produced masks of the whole prostate gland. This dataset is used as an independent test set.

Datasets were split as in Table 7, to define centers characterized by specific image acquisition properties, thus allowing to obtain heterogeneous image distributions among centers. The common preprocessing pipeline applied to all the data comprised of flipping, cropping/padding to the same dimension, and intensity normalization. N4-bias-correction has also been applied to the data from Promise12 to compensate for the intensity artifacts introduced by the endorectal coil. Figure 3 shows an example

Table 7: FedProstate Description of the different centers here considered for the distributed learning scenario, derived by partitioning the four datasets Decathlon, ProstateX, Promise12, and PrivateDS.

ID	#Samples	Dataset	Subset Selection	Splitting Strategy
Local0	32	Decathlon	Full Dataset	5-fold CV
Local1	23	Promise12	No Endorectal Coil	5-fold CV
Local2	27	Promise12	Only Endorectal Coil	5-fold CV
Local3	184	ProstateX	Only Scanner Skyra	5-fold CV
Local4	5	ProstateX	Only Scanner Triotim	External test set
Local5	36	PrivateDS	Full Dataset	External test set

of the resulting splits.

Architectures and Parameters

The segmentation problem for this dataset was addressed by training a 3D UNet architecture with residual connections³¹. The training was based on optimizing the DSC, by using the AdamW optimizer for all experiments³² but for SCAFFOLD, for which we used an SGD optimizer. The UNet implementation is available in the MONAI library². The hyperparameters used for the training phase are available in the Table 8

Table 8: FedProstate Hyperparameters and respective values explored during the tuning phase. Selected value in **bold**. The selection of dropout value was driven by the need to use it for the UBE method. In red, the values selected for FedAdam, FedYogi and FedAdagrad, for which a different tuning was required.

Parameter	Values
Learning Rate	0.0001; 0.001 ; 0.01 ; 0.1; 1
Batch Size	4, 8 , 16
Dropout	0.1, 0.3 ,0.5
Local Steps	10, 15, 20 , 25

5.2 FedHeart

The dataset is a collection of tabular data, and the task consists of binary classification to recognize heart disease. In FLamby¹⁵, a federated version has been proposed by using as subset selection criteria the hospital that provided the data. There are four hospitals (Cleveland’s, Hungarian, Switzerland, and Long Beach), so four clients. The preprocessing in FLamby has been done by removing missing values and encoding non-binary categorical variables as dummy variables.

Architectures and Parameters

A fully connected ReLU network with one hidden layer was used as a classification model. The training was based on optimizing the cross-entropy loss, with an AdamW optimizer. The details on the hyperparameters used for the centralized and local training are available in Table 9.

5.3 FedIXI

The dataset contains T1 and T2 brain MRIs, as well as brain segmentation masks. In FLamby, the federation is obtained by using the hospital as a subset selection criteria, comprising three clients. Pre-processing pipeline comprehends volume resizing to $48 \times 60 \times 48$ voxels, and sample-wise intensity normalization.

²<https://monai.io/index.html>

Table 9: FedHeart Hyperparameters and respective values explored during the tuning phase. Selected value in **bold**. The selection of dropout value was driven by the need to use it for the UBE method.

Parameter	Values
Learning Rate	0.0001; 0.001; 0.01; 0.1 ; 1
Batch Size	8 ; 16; 32; 64
Dropout	0.1; 0.2 ; 0.3
Local Steps	5; 10 ; 20
Centralized epochs	10; 50 ; 100

Architectures and Parameters

The network used as a baseline model for this problem was a 3D UNet³¹. The training was based on optimizing the DICE Loss, by using the AdamW optimizer. The details on the hyperparameters used for the centralized and local training are available in Table 10.

Table 10: FedIXI Hyperparameters and respective values explored during the tuning phase. Selected value in **bold**. The selection of dropout value was driven by the need to use it for the UBE method.

Parameter	Values
Learning Rate	0.0001; 0.001 ; 0.01; 0.1; 1
Batch Size	2 ; 4; 6
Dropout	0.1; 0.2 ; 0.3
Local Steps	5; 10 ; 20
Centralized epochs	10 ; 20; 30

5.4 FedISIC

This dataset represents a skin cancer detection problem through image classification of CT scans. There are 8 different classes, with high distribution imbalance. Starting from the data available in the ISIC2019 dataset³³ the authors of FLamby obtained a federated dataset by allocating to a different client data obtained by using a different scan, so obtaining 6 different clients for the federation. The applied preprocessing is described in³⁴.

Architectures and Parameters

The network used as a baseline model for this problem was EfficientNet²⁵, as in FLamby. The training was based on optimizing the weighted focal loss, by using the AdamW optimizer. The details on the hyperparameters used for the centralized and local training are available in Table 11. When choosing the batch size, we applied

different values for different local clients to account for the high variance in dataset dimensions. The final values are Local 0 and Local 1: 128, Local 2: 64, Local 3, Local 4 and Local 5: 32.

Table 11: FedISIC Hyperparameters and respective values explored during the tuning phase. Selected value in **bold**. The selection of dropout value was driven by the need to use it for the UBE method.

Parameter	Values
Learning Rate	0.0005 ; 0.005 0.001;
Batch Size	32; 64; 128 ; 256
Dropout	0.1; 0.2; 0.3
Local Steps	5; 10 ; 20
Centralized epochs	50; 100; 150

5.5 FedTCGA-BRCA

This dataset is composed of data from the TCGA-GDC portal, specifically those belonging to the breast cancer study, which includes features gathered from 1066 patients. The federated version is obtained by splitting the original dataset into 6 subsets, one for each extraction site, grouped into geographic regions. The task consists of predicting survival outcomes based on the patients’ tabular data, with the event to predict death. Each patient is defined by 38 features.

Architectures and Parameters

As a baseline, we used a fully connected LeakyReLU network. The training was based on optimizing the weighted focal loss, by using the AdamW optimizer. The details on the hyperparameters used for the centralized and local training are available in Table 12.

Table 12: FedTCGA-BRCA Hyperparameters and respective values explored during the tuning phase. Selected value in **bold**. The selection of dropout value was driven by the need to use it for the UBE method.

Parameter	Values
Learning Rate	0.1; 0.01 ; 0.001;
Batch Size	4, 8, 16
Dropout	0.1, 0.2 , 0.3
Local Steps	5, 10 , 15
Centralized epochs	20, 50, 60

5.6 FedKiTS

The KiTS19 dataset is part of the Kidney Tumor Segmentation Challenge 2019 and contains CT scans of 210 patients along with the segmentation masks from 79 hospitals. In FLamby, the federated dataset is defined by splitting scans among different clients based on the providing hospital; they extracted a 6-client federated version by removing hospitals with less than 10 training samples. The preprocessing pipeline comprises intensity clipping followed by intensity normalization and resampling of all the cases to a common voxel spacing.

Architectures and Parameters

As a baseline, following FLamby implementation, a nnUNet³⁵ was used. The training was based on the optimization of the DICE Loss, by using the AdamW optimizer. The details on the hyperparameters used for the centralized and local training are available in Table 13.

Table 13: FedKiTS Hyperparameters and respective values explored during the tuning phase. Selected value in **bold**.

Parameter	Values
Learning Rate	0.1; 0.01 ; 0.001;
Batch Size	2, 4
Dropout	0.1, 0.2 , 0.3
Local Steps	30, 50, 100
Centralized epochs	2000, 5000 , 8000
Optimizer	Adam, AdamW , SGD

5.7 FeTS

The data are gathered from the FeTS 2022 Challenge. A more specific data description is available at the url³. The dataset contains 1251 instances. Following the guidelines, we used a natural partitioning by the institution, obtaining a federation with 23 clients, for which the dataset size is available in Figure 4.


For each patient in the study, the following data modalities are available: native (T1), post-contrast T1-weighted (T1Gd), T2-weighted, and Fluid Attenuated Inversion Recovery (T2-FLAIR) volumes. An example of data for one patient is available in Figure 5.

Data in this FD are very heterogeneous, being acquired with different clinical protocols and various scanners from multiple data-contributing institutions.

The task is multi-class 3D image segmentation, and the labels are the GD-enhancing tumor (ET — label 4), the peritumoral edematous/invaded tissue (ED — label 2), and the necrotic tumor core (NCR — label 1). We segmented single sub-regions, not into intersections as was proposed in the FeTS challenge.

³<https://www.synapse.org/#!/Synapse:syn28546456/wiki/617246>

Fig. 4: Distributions of dataset size among all the clients in the FeTS dataset



Architectures and Parameters

The network used is a SegResNet³⁶, which takes as input the multi-modal data and produces a segmentation mask. The model was trained by optimizing a DICE loss by an AdamW optimizer. As a preprocessing step, each data has been cropped to a common shape of 240x240x128 and intensity normalization has been applied. The details on the hyperparameters used for the centralized and local training are available in Table 14.

Results on local clients

Figure 6 shows the average DSC across splits and runs for the FeTS dataset. This figure completes Figure 2g.

6 Additional results

The Friedman test is a nonparametric method for evaluating the significance of differences between multiple classification algorithms against multiple datasets, comparing how well each model ranks (in terms of accuracy) across different datasets. In Section

Fig. 5: Examples of different modalities for one patient in the FeTS dataset. MASK represents the segmentation mask, which is used as ground truth for our segmentation problem.




Table 14: FeTS Hyperparameters and respective values explored during the tuning phase. Selected value in **bold**.

Parameter	Values
Learning Rate	0.1; 0.01; 0.001 ;
Batch Size	1, 2, 4
Dropout	0.1, 0.2, 0.3
Local Steps	1, 6 , 20
Centralized epochs	15, 30 , 50

3.1 we have shown the numerical results of the Friedman test, while in Figure 7 we show an analysis of the rankings on the various methods. We can read the figure as the representation of the probability of each method of ranking at a given position when tested on a dataset. The sparsity of the heatmap qualitatively suggests that no method emerges above the others to be systematically the best. These results are consistent with those of the statistical analysis of the p-values.

References

- [1] Sheller MJ, Edwards B, Reina GA, Martin J, Pati S, Kotrotsou A, et al. Federated learning in medicine: facilitating multi-institutional collaborations without


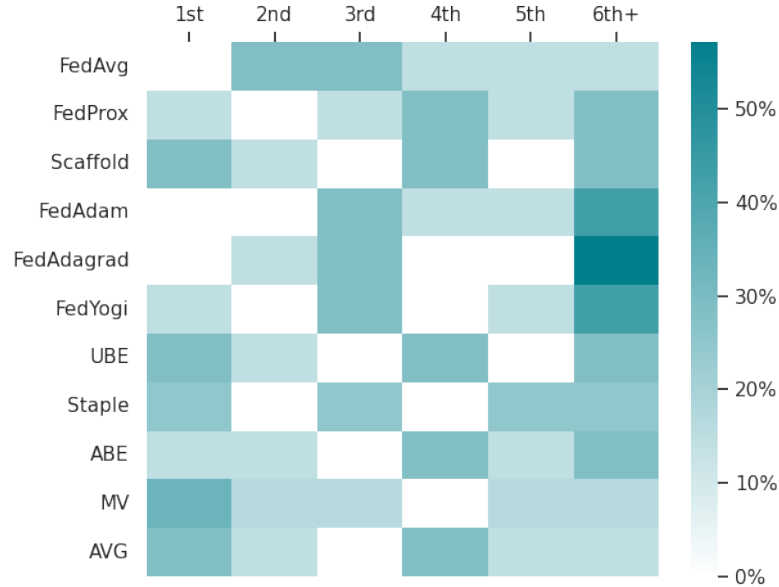


Fig. 6: Average dice score obtained by locally-trained models on the FeTS dataset

sharing patient data. *Scientific reports*. 2020;10(1):12598.

- [2] Rieke N, Hancox J, Li W, Milletari F, Roth HR, Albarqouni S, et al. The future of digital health with federated learning. *NPJ digital medicine*. 2020;3(1):119.
- [3] Collaborative learning without sharing data. *Nature Machine Intelligence*. 2021;3(459).
- [4] Usynin D, Ziller A, Makowski M, Braren R, Rueckert D, Glocker B, et al. Adversarial interference and its mitigations in privacy-preserving collaborative machine learning. *Nature Machine Intelligence*. 2021;3(9):749–758.
- [5] Kairouz P, McMahan HB, Avent B, Bellet A, Bennis M, Bhagoji AN, et al. Advances and open problems in federated learning. *Foundations and Trends® in Machine Learning*. 2021;14(1–2):1–210.
- [6] Joshi M, Pal A, Sankarasubbu M. Federated learning for healthcare domain-pipeline, applications and challenges. *ACM Transactions on Computing for Healthcare*. 2022;3(4):1–36.
- [7] Guha N, Talwalkar A, Smith V. One-shot federated learning. *arXiv preprint arXiv:190211175*. 2019;.
- [8] Wang H, Yurochkin M, Sun Y, Papailiopoulos D, Khazaeni Y. Federated learning with matched averaging. *arXiv preprint arXiv:200206440*. 2020;.

Fig. 7: Performance ranking. Columns represent the ranking position for the accuracy of the methods across datasets. We note that not all the methods can be applied to each dataset: STAPLE can only be applied to segmentation tasks, and MV does not apply to the survival task of FedTGCA-BRCA. Overall, no method stands out in terms of overall best performance (Friedman test on segmentation ($p = 0.73$), classification ($p = 0.62$), and survival analysis ($p = 0.42$) tasks).



- [9] Innocenti L, Antonelli M, Cremonesi F, Sarhan K, Granados A, Goh V, et al. Benchmarking Collaborative Learning Methods Cost-Effectiveness for Prostate Segmentation. arXiv preprint arXiv:230917097. 2023;.
- [10] Lee GH, Shin SY. Federated learning on clinical benchmark data: performance assessment. Journal of medical Internet research. 2020;22(10):e20891.
- [11] Chen D, Gao D, Kuang W, Li Y, Ding B. pFL-bench: A comprehensive benchmark for personalized federated learning. Advances in Neural Information Processing Systems. 2022;35:9344–9360.
- [12] Caldas S, Duddu SMK, Wu P, Li T, Konečný J, McMahan HB, et al. Leaf: A benchmark for federated settings. arXiv preprint arXiv:181201097. 2018;.
- [13] Chaudhari H, Jagielski M, Oprea A. SafeNet: The Unreasonable Effectiveness of Ensembles in Private Collaborative Learning. In: 2023 IEEE Conference on Secure and Trustworthy Machine Learning (SaTML). IEEE; 2023. p. 176–196.

- [14] Gupta S, Kumar S, Chang K, Lu C, Singh P, Kalpathy-Cramer J. Collaborative privacy-preserving approaches for distributed deep learning using multi-institutional data. *RadioGraphics*. 2023;43(4):e220107.
- [15] du Terrail JOd, Ayed SS, Cyffers E, Grimberg F, He C, Loeb R, et al. Flamby: Datasets and benchmarks for cross-silo federated learning in realistic healthcare settings. *arXiv preprint arXiv:221004620*. 2022;.
- [16] Audibert J, Michiardi P, Guyard F, Marti S, Zuluaga MA. Usad: Unsupervised anomaly detection on multivariate time series. In: *Proceedings of the 26th ACM SIGKDD international conference on knowledge discovery & data mining*; 2020. p. 3395–3404.
- [17] McMahan B, Moore E, Ramage D, Hampson S, y Arcas BA. Communication-Efficient Learning of Deep Networks from Decentralized Data. In: *ICML 2017*; 2017. .
- [18] Li T, Sahu AK, Zaheer M, Sanjabi M, Talwalkar A, Smith V. Federated Optimization in Heterogeneous Networks. *Proceedings of the 1 st Adaptive & Multitask Learning Workshop, Long Beach, California, 2019*. 2018;p. 1–28. [arXiv:1812.06127](https://arxiv.org/abs/1812.06127).
- [19] Karimireddy SP, Kale S, Mohri M, Reddi S, Stich S, Suresh AT. Scaffold: Stochastic controlled averaging for federated learning. In: *International conference on machine learning*. PMLR; 2020. p. 5132–5143.
- [20] Reddi S, Charles Z, Zaheer M, Garrett Z, Rush K, Konečný J, et al. Adaptive federated optimization. *arXiv preprint arXiv:200300295*. 2020;.
- [21] Safdar K, Akbar S, Shoukat A. A majority voting based ensemble approach of deep learning classifiers for automated melanoma detection. In: *2021 International Conference on Innovative Computing (ICIC)*. IEEE; 2021. p. 1–6.
- [22] Warfield SK, Zou KH, Wells WM. Simultaneous truth and performance level estimation (STAPLE): an algorithm for the validation of image segmentation. *IEEE transactions on medical imaging*. 2004;23(7):903–921.
- [23] Ruta D, Gabrys B. An overview of classifier fusion methods. *Computing and Information systems*. 2000;7(1):1–10.
- [24] Adiga A, Wang L, Hurt B, Peddireddy A, Porebski P, Venkatramanan S, et al. All models are useful: Bayesian ensembling for robust high resolution covid-19 forecasting. In: *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*; 2021. p. 2505–2513.
- [25] Tan M, Le Q. Efficientnet: Rethinking model scaling for convolutional neural networks. In: *International conference on machine learning*. PMLR; 2019. p.

6105–6114.

- [26] Pati S, Baid U, Zenk M, Edwards B, Sheller M, Reina GA, et al. The federated tumor segmentation (fets) challenge. arXiv preprint arXiv:210505874. 2021;.
- [27] Antonelli M, Reinke A, Bakas S, Farahani K, Kopp-Schneider A, Landman BA, et al. The medical segmentation decathlon. *Nature communications*. 2022;13(1):4128.
- [28] Litjens G, Toth R, Van De Ven W, Hoeks C, Kerkstra S, Van Ginneken B, et al. Evaluation of prostate segmentation algorithms for MRI: the PROMISE12 challenge. *Medical image analysis*. 2014;18(2):359–373.
- [29] Armato III SG, Huisman H, Drukker K, Hadjiiski L, Kirby JS, Petrick N, et al. PROSTATEx Challenges for computerized classification of prostate lesions from multiparametric magnetic resonance images. *Journal of Medical Imaging*. 2018;5(4):044501–044501.
- [30] Cuocolo R, Stanzione A, Castaldo A, De Lucia DR, Imbriaco M. Quality control and whole-gland, zonal and lesion annotations for the PROSTATEx challenge public dataset. *European Journal of Radiology*. 2021;138:109647.
- [31] Ronneberger O, Fischer P, Brox T. U-net: Convolutional networks for biomedical image segmentation. In: *MICCAI 2015*. Springer; 2015. .
- [32] Loshchilov I, Hutter F. Decoupled weight decay regularization. In: arXiv preprint arXiv:1711.05101; 2017. .
- [33] Gutman D, Codella NC, Celebi E, Helba B, Marchetti M, Mishra N, et al. Skin lesion analysis toward melanoma detection: A challenge at the international symposium on biomedical imaging (ISBI) 2016, hosted by the international skin imaging collaboration (ISIC). arXiv preprint arXiv:160501397. 2016;.
- [34] Arora A.: Siim-istic melanoma classification - my journey to a top 5 Available from: <https://amaarora.github.io/2020/08/23/siimistic.html>.
- [35] Isensee F, Jaeger PF, Kohl SA, Petersen J, Maier-Hein KH. nnU-Net: a self-configuring method for deep learning-based biomedical image segmentation. *Nature methods*. 2021;18(2):203–211.
- [36] Myronenko A. 3D MRI brain tumor segmentation using autoencoder regularization. *Brainlesion Glioma Mult. Scler. Stroke Trauma Brain Inj-BrainLes*. 2018;2019:11384.