

Fed-Sophia: A Communication-Efficient Second-Order Federated Learning Algorithm

Ahmed Elbakary^{1,2}, Chaouki Ben Issaid¹, Mohammad Shehab¹, Karim Seddik², Tamer ElBatt^{3,4}, Mehdi Bennis¹

¹Centre for Wireless Communications (CWC), University of Oulu, Finland

²Department of Electronics and Communications Engineering, American University in Cairo, Egypt

³Department of Computer Science and Engineering, American University in Cairo, Egypt

⁴Department of Electronics and Communications Engineering, Cairo University, Egypt

Abstract—Federated learning is a machine learning approach where multiple devices collaboratively learn with the help of a parameter server by sharing only their local updates. While gradient-based optimization techniques are widely adopted in this domain, the curvature information that second-order methods exhibit is crucial to guide and speed up the convergence. This paper introduces a scalable second-order method, allowing the adoption of curvature information in federated large models. Our method, coined Fed-Sophia, combines a weighted moving average of the gradient with a clipping operation to find the descent direction. In addition to that, a lightweight estimation of the Hessian’s diagonal is used to incorporate the curvature information. Numerical evaluation shows the superiority, robustness, and scalability of the proposed Fed-Sophia scheme compared to first and second-order baselines.

Index Terms—Distributed optimization, federated learning, second-order methods, deep neural networks, communication efficiency.

I. INTRODUCTION

Recent advances in edge computing have paved the way to enable new machine learning paradigms (ML). One promising paradigm is federated learning (FL) [1], [2], where devices preserve their local data and only share, with a parameter server (PS), their local models. Privacy preservation and data locality are two major benefits of FL, in contrast to traditional centralized ML models, which require the entire dataset to be stored in a central cloud. On the contrary, FL relies on a two-stage approach. First, each device performs local training independently, using its own local data. After one or more iterations of local training, each device sends its local model or gradient to the PS. Afterward, the models are aggregated using one of the known aggregation schemes in the literature [3]. Finally, after aggregation, the PS shares the refined model with all devices for enhanced performance. This process continues in an iterative manner over multiple communication rounds until convergence. The ultimate goal of FL is to reach a consensus and find a global model that solves the task at hand.

When using first-order methods, clients employ the gradient to find a descent direction that could guide the convergence. Meanwhile, the default second-order optimization technique is Newton’s method, which requires computing the full Hessian and using it as a pre-conditioner for the gradient. In our work, instead of computing the full Hessian as in Newton’s

method, which comes with a heavy computational cost, we rely on an estimation of the Hessian’s diagonal to incorporate the curvature information across each dimension into the convergence process. One problem stemming from Newton’s method is that it might converge to a saddle point instead of a local minimum. Another issue with Hessian-based methods is that the Hessian entries might mislead the convergence process by capturing non-accurate curvature information due to sudden changes in the loss function’s curvature [4]. To mitigate these problems, a clipping operation is introduced to guard against non-positive entries of the Hessian.

The main contributions of this paper can be summarized as follows

- We propose Federated Sophia (Fed-Sophia), a scalable second-order method that enables large models to make use of the curvature information of the loss function. In addition to that, Fed-Sophia can handle heterogeneous settings where the data distribution differs among devices.
- The proposed method uses a lightweight estimation of the Hessian and only shares the model’s parameter vector, making it computation and communication-efficient.
- Several numerical experiments are conducted to show that Fed-Sophia outperforms other baselines for solving an image classification task using convolutional neural network (CNN) or multi-layer perception (MLP) models. Furthermore, we compare the energy consumption and the carbon footprint of Fed-Sophia to other baselines.

The rest of this paper is organized as follows. We highlight some of the available methods in Section II. In Section III, the problem formulation and the system model are introduced. We introduce our proposed method in Section IV. The details of our experimental results are laid out in Section V, where the proposed method is compared to other baselines.

II. RELATED WORK

In what follows, we highlight some differences between first and second-order methods and motivate the need for scalable optimizers in the FL setting.

A. First-order FL Methods

Federated averaging (FedAvg) introduced a generalization of stochastic gradient descent (SGD). The underlying idea is to

have every device run several local updates before sending the local model parameters to the PS. The server then averages the local models' parameters and sends back the refined, global model parameters. In [5], the authors proposed a federated version of a set of popular centralized optimization methods like AdaGrad [6] and Adam [7], using the same two-stage approach employed in FedAvg. To handle data heterogeneity, Scaffold [8] used variance reduction to overcome the drift in local updates at each device. In [9], the authors reduced the communication overhead and the data heterogeneity using both compression and gradient tracking. The fundamental problem with first-order methods is their inability to capture the curvature information of the loss landscape, which leads to a uniform step size across all dimensions of the loss function.

B. Second-order FL Methods

Second-order methods have the advantage of capturing the effect of the curvature information of the loss function in the training process. FedNL [10] is the first framework for using Newton's method in FL settings. The paper introduced compression techniques to improve communication efficiency between the PS and the devices. Still, the problem of computing the full Hessian at each device persists, making the idea impractical for large models. FedNew [11] added another layer of privacy by hiding the gradient information. It also proposed stochastic quantization as a solution for the problem of communicating the Hessian with the PS while still approximating the inverse-Hessian-gradient product. FLECS [12] proposed a framework for second-order methods, utilizing a lower dimensional representation of the Hessian to make the computation lightweight. A major bottleneck in this idea is that the server has to save the full Hessian. Meanwhile, DONE [13] uses an approximate Newton-type method and employs Richardson's iteration to find an approximation for the search direction at each client. Furthermore, DONE requires a large number of local iterations to converge, making it unreliable for practical settings.

C. Optimizers in Deep Learning

The choice of an optimizer in deep learning is often a task-dependent case. For example, SGD has dominated the computer vision field, while Adam [7] excels in transformers-based language models. On the other hand, second-order methods are not widely adopted in training deep neural networks (DNN). One attempt to apply second-order methods is AdaHessian [14], a second-order method that estimates the Hessian's diagonal using Hutchinson's method [15] with a spatial averaging mechanism to control misleading Hessian entries. In [16], the authors proposed Sophia, a second-order optimizer that exploits a lightweight estimation for the diagonal of the Hessian to incorporate curvature information into the optimization process. Originally introduced as an optimizer for large language models (LLMs), one advantage of Sophia is that it does not estimate the diagonal of the Hessian in each iteration, a key feature that can be well leveraged in FL settings.

III. SYSTEM MODEL AND PROBLEM FORMULATION

Consider a network of N devices and a PS, where each device can communicate with the PS only. Each device has its local data set $\mathcal{D}_i(\mathbf{X}, \mathbf{y})$ and loss function $f_i(\boldsymbol{\theta}) : \mathbb{R}^d \rightarrow \mathbb{R}$, where \mathbf{X} is the features matrix, \mathbf{y} is the label vector, and $\boldsymbol{\theta}$ is the ML model we aim to train, e.g., the weights and biases of a NN. Devices train a local model $\boldsymbol{\theta}_i \in \mathbb{R}^d$ and send the trained model to the PS. The server can aggregate the information sent from the devices using some aggregation scheme to guide the convergence process. After aggregation, the server sends back the updated information to all clients, which starts another round of local training. This process continues until convergence. The objective of the training process is to perform empirical risk minimization in the form

$$\min_{\boldsymbol{\theta} \in \mathbb{R}^d} f(\boldsymbol{\theta}) \triangleq \frac{1}{N} \sum_{i=1}^N f_i(\boldsymbol{\theta}). \quad (1)$$

In the remainder of this paper, let $\mathbf{H}_i \triangleq \nabla^2 f_i(\boldsymbol{\theta})$ and $\mathbf{g}_i \triangleq \nabla f_i(\boldsymbol{\theta})$ denote the Hessian and the gradient of the loss function of the i th device, respectively. The entries of the Hessian matrix could be divided into two classes: the diagonal elements and the off-diagonal elements. The diagonal elements indicate the curvature information across each dimension. High values for diagonal elements mean steep curvature, which requires smaller steps. Meanwhile, small values reflect a flat dimension, in which case faster steps could be adopted. On the other hand, off-diagonal elements represent the cross-curvature information. To illustrate the effect of the Hessian's elements, let's consider the following loss function

$$f(\theta_1, \theta_2) = \theta_1^2 + 2\theta_1\theta_2 + 3\theta_2^2, \quad (2)$$

where $\boldsymbol{\theta} = [\theta_1, \theta_2]$ is the model we aim to learn. The Hessian matrix for this function is given by

$$\mathbf{H} = \nabla^2 f(\boldsymbol{\theta}) = \begin{bmatrix} \frac{\partial^2 f}{\partial \theta_1^2} & \frac{\partial^2 f}{\partial \theta_1 \partial \theta_2} \\ \frac{\partial^2 f}{\partial \theta_2 \partial \theta_1} & \frac{\partial^2 f}{\partial \theta_2^2} \end{bmatrix} = \begin{bmatrix} 2 & 2 \\ 2 & 6 \end{bmatrix}. \quad (3)$$

The optimal step size to take across a dimension should be proportional to the Hessian entries corresponding to those dimensions. For example, the optimal step size for θ_1 and θ_2 should be proportional to $1/2$ and $1/6$, respectively. Those entries are the diagonal of the Hessian. Figure 1 demonstrates the effect of the optimal step size on the convergence. As we can see, the gradient-based method using uniform step size takes a lot of steps to converge, while a Hessian-based method can converge in two steps.

IV. PROPOSED ALGORITHM

Our proposed algorithm, Fed-Sophia, is a second-order method that leverages the Sophia optimizer [16] into the FL setting. The algorithm's computational and communication efficiency is attributed to two primary factors: (i) the lightweight estimation of the Hessian matrix, which, owing to its diagonal structure, is analogous to computing a vector, and (ii) the Hessian estimation does not occur in every communication

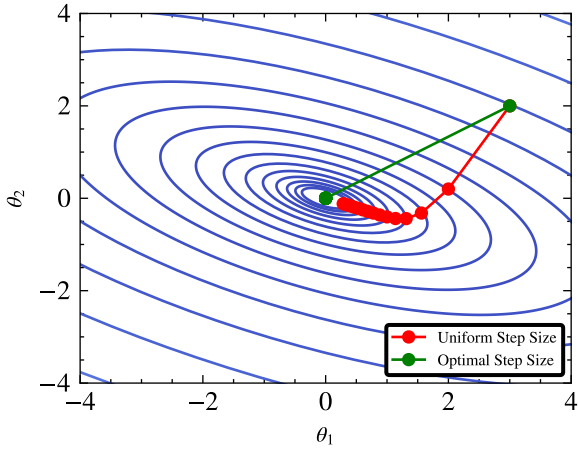


Fig. 1: The effect of the Hessian-based step size against the gradient-based method.

round. Similar to FedAvg, every device in Fed-Sophia shares its model's parameters with the PS. Then, the server aggregates the parameters using their average as follows

$$\Theta^{k+1} = \frac{1}{N} \sum_{i=1}^N \theta_i^k, \quad (4)$$

where Θ is the server's parameters at iteration $(k+1)$, and θ_i is the model's parameters of the i th device. As shown in [1], [17], this idea of averaging the models works surprisingly well with over-parameterized NNs.

A. Hessian as a Pre-conditioner

One problem with first-order methods, like FedAvg, is their inability to incorporate curvature information during the optimization process. Curvature information is best obtained using the Hessian matrix. At the k th communication round, a general descent direction can be obtained as follows

$$\Delta \theta_i^k \triangleq \theta_i^{k+1} - \theta_i^k = (\mathbf{H}_i^k)^{-s} \mathbf{g}_i^k. \quad (5)$$

Herein, $s = 1$ corresponds to the vanilla Newton method, while $s = 0$ rolls back to the gradient descent method. The idea of the Hessian as a pre-conditioner, as shown in (5), is that it adapts the step size taken along the descent direction by $(\mathbf{H}_i^k)^{-1}$.

Computing the full Hessian is almost impossible in the case of a very large NN. This makes certain FL algorithms impractical in real-world settings. One way to estimate the Hessian is by computing the diagonal information only [18]. The diagonal elements of the Hessian contain curvature information across each dimension. Thus, (5) can be approximated as

$$\Delta \theta_i^k \approx \text{diag}(\mathbf{H}_i^k)^{-1} \mathbf{g}_i^k. \quad (6)$$

Several techniques for such diagonal estimation have been proposed, e.g., Hutchinson's and Gauss-Newton-Bartlett (GNB) methods [16]. Fed-Sophia uses the GNB estimator to estimate the Hessian as it has been shown to exhibit a better generalization performance [16]. Given a loss function $f(\theta)$, the

Algorithm 1: Federated Sophia (Fed-Sophia)

- 1: **Parameters:** learning rate η , number of local iterations J , hyperparameters $(\lambda, \beta_1, \beta_2, \epsilon, \rho, \tau)$.
 - 2: **Initialize:** $\Theta^0, m^0 = \mathbf{0}, h_i^{-\tau} = \mathbf{0}$
 - 3: **for** communication round $k = 0$ to K **do**
 - 4: **for** each device $i \in [N]$ **do**
 - 5: Receive the global model Θ^k and set the local model as $\theta_i^k = \Theta^k$
 - 6: **for** local iterations $j = 0$ to J **do**
 - 7: Compute the local gradient $\mathbf{g}_i^k = \nabla h_i(\theta_i^k)$
 - 8: Update the local moving average of the gradient: $m_i^k = \beta_1 m_i^{k-1} + (1 - \beta_1) \mathbf{g}_i^k$
 - 9: **if** $k \bmod \tau = 0$ **then**
 - 10: Compute the local estimated Hessian: $\hat{h}_i^k = \text{Estimator}(\theta_i^k)$ using GNB Algorithm 2
 - 11: Update the local moving average of the Hessian estimator: $h_i^k = \beta_2 h_i^{k-\tau} + (1 - \beta_2) \hat{h}_i^k$
 - 12: **else**
 - 13: $h_i^k = h_i^{k-1}$
 - 14: **end if**
 - 15: Update the local model with weight decay: $\theta_i^k = \theta_i^k - \eta \lambda \theta_i^k$
 - 16: Update the local model: $\theta_i^{k+1} = \theta_i^k - \eta \text{clip}\left(\frac{m_i^k}{\max(h_i^k, \epsilon)}, \rho\right)$
 - 17: **end for**
 - 18: Send the updated local models θ_i^{k+1} back to the server
 - 19: **end for**
 - 20: Update the global model: $\Theta^{k+1} = \frac{1}{N} \sum_{i=1}^N \theta_i^{k+1}$
 - 21: **end for**
-

Hessian matrix can be computed using the following Gauss-Newton decomposition

$$\mathbf{H} = \mathbf{J}_\theta \phi(\theta, \mathbf{x}) \cdot \mathbf{S} \cdot \mathbf{J}_\theta \phi(\theta, \mathbf{x})^T + \mathbf{J}_{\theta\theta} \phi(\theta, \mathbf{x}) \cdot \mathbf{q}, \quad (7)$$

where \mathbf{J} is the Jacobian of ϕ with respect to θ , ϕ is the logits function or the mapping from input to raw output values, \mathbf{S} is the second-order derivative of the loss with respect to the logits, $\mathbf{J}_{\theta\theta} \phi(\theta, \mathbf{x})$ is the second-order derivatives of the logits function with respect to θ , and \mathbf{q} is the first-order derivative of the loss with respect to the logits. The first term is often called the Gauss-Newton matrix, while the second is usually very small and could be negligible. The idea is to estimate the diagonal of the Gauss-Newton matrix and use this as the Hessian's diagonal. The procedure to estimate the diagonal of the Hessian is depicted in Algorithm 2, where \odot denotes the element-wise product. Hence, the estimated Hessian for the k th communication round is given by

$$\hat{h}_i^k = \text{Estimator}(\theta_i^k) \quad (8)$$

where the estimator is the GNB method in our case.

A big advantage of Fed-Sophia is its ability to compute the estimated Hessian's diagonal every τ iteration, where τ is an

Algorithm 2: Gauss-Newton-Bartlett (GNB)

- 1: **Parameters:** θ
 - 2: Draw a mini-batch of the input $\{\mathbf{x}_b\}_{b=1}^B$
 - 3: Compute the logits on the mini-batch $\{\phi(\theta, \mathbf{x}_b)\}_{b=1}^B$
 - 4: Sample $\hat{\mathbf{y}}_b \sim \text{Softmax}(\phi(\theta, \mathbf{x}_b))$
 - 5: Calculate $\hat{\mathbf{g}} = \nabla(1/B \sum f(\phi(\theta, \mathbf{x}_b), \hat{\mathbf{y}}_b))$
 - 6: return $B \cdot \hat{\mathbf{g}} \odot \hat{\mathbf{g}}$
-

integer usually between 1 and 10. Another advantage for Fed-Sophia is that it uses mini-batches to compute the loss, not the complete data. This might reflect some noise in the gradient and the Hessian. To avoid this issue and smoothen both the gradient and the Hessian, an exponential moving average is computed instead of the raw value for both of them. The update for the gradient is given by

$$\mathbf{m}_i^k = \beta_1 \mathbf{m}_i^{k-1} + (1 - \beta_1) \mathbf{g}_i^k, \quad (9)$$

while the Hessian is updated, every τ iterations, using

$$\mathbf{h}_i^k = \beta_2 \mathbf{h}_i^{k-\tau} + (1 - \beta_2) \hat{\mathbf{h}}_i^k, \quad (10)$$

where β_1 and β_2 are two hyperparameters.

B. Adaptive Step Size

The update at the i th device is performed by dividing the moving average of the gradient by the moving average of the estimated diagonal Hessian, i.e., $\mathbf{m}_i^k / \mathbf{h}_i^k$. The motivation for this update is the fact that the ideal step size to take should be proportionally inverse to the Hessian across each dimension to avoid uniform steps across all dimensions. In other words, we need to make weighted steps at each dimension, where the weights are the Hessian's information across those dimensions. Rather than using a fixed update step across all dimensions of the loss function, the idea is to adjust this step to enable sharp dimensions and flat dimensions to behave accordingly. For sharp dimensions, the step size is expected to be less than in flat dimensions. Dividing by the estimated Hessian enables the adaptive step size mechanism, which in turn assigns different step sizes depending on the shape of the loss curvature.

C. Clipping Operation

Along with the problem of converging to a saddle point in the case of a non-convex loss function, the Hessian entries might be misleading the convergence process in case there are sudden changes along the loss function curvature. To help mitigate these problems, a clipping operation is introduced. The clipping operation guards against inaccurate information that might be incorporated into the Hessian's diagonal. Given a vector \mathbf{z} , the formula for the clipping operation is given by

$$\text{clip}(\mathbf{z}, \rho) = \max\{\min\{z, \rho\}, -\rho\}, \quad (11)$$

where ρ is a positive real number that controls the maximum update magnitude. Using the clipping operation, the descent direction (6) becomes

$$\Delta \theta_i^k = \text{clip} \left(\frac{\mathbf{m}_i^k}{\max\{\mathbf{h}_i^k, \epsilon\}}, \rho \right), \quad (12)$$

where \mathbf{m}_i^k and \mathbf{h}_i^k are computed according to (9) and (10), respectively, and ϵ is a very small positive constant used to prevent division by zero. Algorithm 1 shows the full details of Fed-Sophia.

V. NUMERICAL RESULTS

In this section, we lay out the details of the experiments along with the results. First, we compare the performance of Fed-Sophia to the other baselines in terms of communication and computation efficiency. Then, we study the effect of the hyperparameters on Fed-Sophia's performance. Finally, we investigate the energy and carbon footprint of all algorithms.

A. Experimental Settings

We conducted several experiments to compare the performance of Fed-Sophia with two baselines: FedAvg, a first-order method that utilizes only the gradient information, and DONE, an approximate Newton-type method. Two data sets are used: MNIST and Fashion MNIST. The data is distributed among 32 devices, and each partition is split into 75% and 25% for training and testing, respectively. All the experiments are in the non-IID setting. In our experiments, we consider two different classification models, MLP and CNN, with the loss function being cross-entropy in both of them. The number of local iterations for both Fed-Sophia and FedAvg is taken to be $J = 10$. For DONE, we tune the number of local iterations to maintain a fair comparison.

To monitor the energy footprint of each device i , both computation and communication costs are computed as follows

$$E_{total}(k) = E_c(k) + E_t(k), \quad (13)$$

with

$$E_c(k) = \sum_{n=1}^k \sum_{j=1}^J e_i^{j,n} \text{ and } E_t(k) = \sum_{n=1}^k b(\theta_i^k) e_{i,PS}^k, \quad (14)$$

where k is the number of communication rounds, $e_i^{j,n}$ is the energy consumed by the device i for one local iteration, $b(\theta_i^k)$ is the size of the model vector in bits, and $e_{i,PS}$ is the energy needed to transmit information between the device i and the PS. The transmitted model's parameters are assumed to be in the 32-bit form. We use the same channel model as in [19] and assume a uniform distance between the server and all clients within a space of size $100 \times 100 m^2$. According to Shannon's formula, the maximum achievable rate for each client is given by $R = B \log_2(1 + \frac{P_t}{d_{i,PS} B N_0})$, where B is the bandwidth,

TABLE I: Effect of the learning rate and the number of local iterations on the test accuracy for the Fashion MNIST dataset with CNN.

Learning rate (η)	Local iterations (J)	Test accuracy(%)
0.01		76.3
0.003	10	80.3
0.0005		71.9
	1	58.7
0.001	5	73.3
	10	76.5

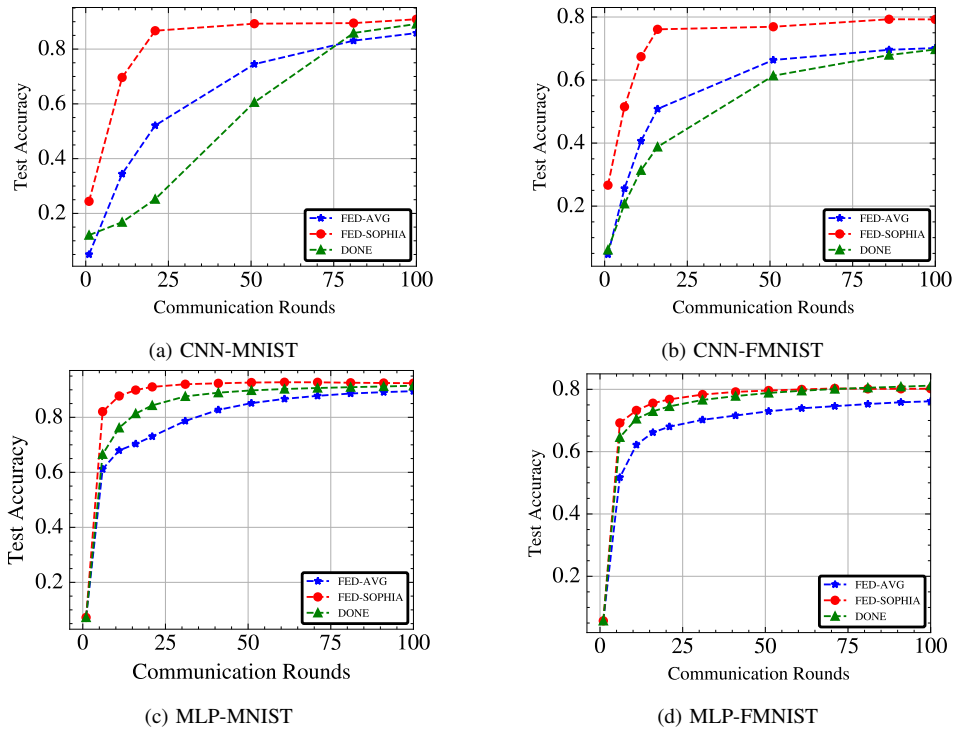


Fig. 2: Test accuracy for Fed-Sophia against other baselines in terms of communication rounds for MNIST/FMNIST datasets using MLP and CNN models.

P_t is the transmission power, $d_{i,PS}$ is the distance between client i and the PS, and N_0 is the noise spectral density. We set $P_t = 100\text{mW}$, $B = 2\text{MHz}$, and $N_0 = 10^{-9}\text{W/Hz}$.

B. Performance Comparison

Herein, we compare the performance of Fed-Sophia against FedAvg and DONE in terms of the test accuracy on both MNIST and FMNIST datasets. For Fed-Sophia and FedAvg, we utilize mini-batches of size 512, a common way to train DNNs. On the other hand, DONE requires the full data for each client, which comes at a computational cost. As shown in Fig. 2, the proposed Fed-Sophia scheme requires fewer communication rounds than FedAvg and DONE in all experiments. It also outperforms both of them in terms of the test accuracy for the CNN model Fig. 2(a) and 2(b) with the two datasets. For the MLP model, Fed-Sophia outperforms the two baselines with the MNIST dataset, requiring only 30 communication rounds to converge compared to 70 and 100 for DONE and FedAvg, respectively as shown in Fig. 2(c). For the FMNIST dataset in Fig. 2(d), DONE and Fed-Sophia are almost identical in terms of test accuracy, with the advantage of early convergence for Fed-Sophia. In Fig. 3, we plot the test accuracy of all algorithms against the total number of iterations to grasp an idea of the computational cost. For the MNIST dataset, it turns out that Fed-Sophia reaches a target test accuracy of 75% after almost 30 iterations, while FedAvg and DONE require 200 and 750 iterations to achieve the same target accuracy, respectively. On the other hand, Fed-Sophia with FMNIST dataset requires 200 iterations to

converge, while FedAvg and DONE require almost the full number of iterations.

C. Effect of the Hyperparameters

In this subsection, we study the effect of the learning rate (η) and the number of local iterations (J) on the performance of Fed-Sophia. We fix $J = 10$ and experiment with different learning rates for the Fashion MNIST dataset. The test accuracy for the three different values of the learning rate is summarized in Table I. Next, we study the effect of the local iterations to see if it is worth adding more computational overhead to achieve some performance gain. As shown in Table I, as we increase the local number of iterations, Fed-Sophia gains a reasonable performance. It is also clear that the performance gain between 5 and 10 local iterations is minimal.

D. Energy-Efficiency and Carbon Footprint

As shown in Table II, Fed-Sophia outperforms DONE and FedAvg in terms of computation energy consumption for a target test accuracy of 75% using a CNN model for the MNIST dataset. Furthermore, DONE consumes up to $4000\times$ computational power than Fed-Sophia. Fed-Sophia requires almost half of the computational energy needed for FedAvg. On the other hand, the communication energy for Fed-Sophia is less than the other two baselines by a large factor. Fed-Sophia only requires almost 20% of the needed communication energy for the closest baseline, i.e., FedAvg. Therefore, the total footprint for Fed-Sophia is minimal, which proves the communication and computation efficiency of our proposed method compared to the other two baselines.

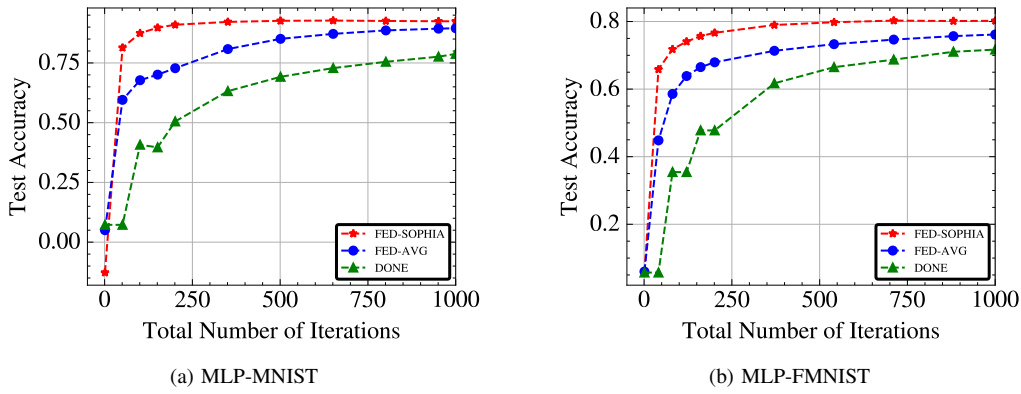








Fig. 3: Test accuracy for Fed-Sophia against other baselines in terms of the number of total iterations for MNIST and FMNIST datasets using MLP.

TABLE II: Computation/communication energy costs and corresponding carbon footprints for MNIST with CNN for target accuracy of 75%.

Algorithm	Computation Energy [MJ]	Communication Energy [MJ]	Total Footprint [kg-CO ₂ -eq]
DONE	$6170E-4$ 	153.2 	$1097E+3$
FedAvg	$2.2E-4$ 	139.2 	$0.123E+3$
Fed-Sophia	$1.4E-4$ 	27.8 	$0.004E+3$

VI. CONCLUSION

In this paper, we presented Fed-Sophia, a scalable and second-order method for FL. Fed-Sophia is a computation and communication-efficient method that leverages the curvature information of the loss function by estimating the diagonal of the Hessian. The update is the moving average of the gradient scaled by the inverse of the estimated diagonal Hessian. Furthermore, Fed-Sophia mitigates the effect of inaccurate entries of the Hessian by augmenting the update into a clipping operation, guarding against misleading curvature information. Our experiments showed that the proposed method outperforms other baselines on image classification tasks.

ACKNOWLEDGMENTS

This work is partially supported by the European Commission through Grant no. 101095363 (Horizon Europe SNS JU ADROIT6G project), Grant no. 101139266 (6G-INTENSE project), and the Academy of Finland, 6G Flagship program (Grant no. 346208).

REFERENCES

- [1] B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. y Arcas, "Communication-efficient learning of deep networks from decentralized data," in *Artificial intelligence and statistics*. PMLR, 2017, pp. 1273–1282.
- [2] P. Kairouz *et al.*, "Advances and open problems in federated learning," *Foundations and Trends® in Machine Learning*, vol. 14, no. 1–2, pp. 1–210, 2021.
- [3] P. Qi, D. Chiaro, A. Guzzo, M. Ianni, G. Fortino, and F. Piccialli, "Model aggregation techniques in federated learning: A comprehensive survey," *Future Generation Computer Systems*, 2023.
- [4] S. P. Boyd and L. Vandenberghe, *Convex optimization*. Cambridge University Press, 2004.
- [5] S. Reddi, Z. Charles, M. Zaheer, Z. Garrett, K. Rush, J. Konečný, S. Kumar, and H. B. McMahan, "Adaptive federated optimization," *arXiv preprint arXiv:2003.00295*, 2020.
- [6] J. Duchi, E. Hazan, and Y. Singer, "Adaptive subgradient methods for online learning and stochastic optimization," *Journal of machine learning research*, vol. 12, no. 7, 2011.
- [7] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.
- [8] S. P. Karimireddy, S. Kale, M. Mohri, S. Reddi, S. Stich, and A. T. Suresh, "Scaffold: Stochastic controlled averaging for federated learning," in *International conference on machine learning*. PMLR, 2020, pp. 5132–5143.
- [9] F. Haddadpour, M. M. Kamani, A. Mokhtari, and M. Mahdavi, "Federated learning with compression: Unified analysis and sharp guarantees," in *International Conference on Artificial Intelligence and Statistics*. PMLR, 2021, pp. 2350–2358.
- [10] M. Safaryan, R. Islamov, X. Qian, and P. Richtarik, "FedNL: Making Newton-type methods applicable to federated learning," in *International Conference on Machine Learning*. PMLR, 2022, pp. 18959–19010.
- [11] A. Elgabli, C. B. Issaid, A. S. Bedi, K. Rajawat, M. Bennis, and V. Aggarwal, "FedNew: A communication-efficient and privacy-preserving Newton-type method for federated learning," in *International Conference on Machine Learning*. PMLR, 2022, pp. 5861–5877.
- [12] A. Agafonov, D. Kamzolov, R. Tappenden, A. Gasnikov, and M. Takáč, "Flecs: A federated learning second-order framework via compression and sketching," *arXiv preprint arXiv:2206.02009*, 2022.
- [13] C. T. Dinh, N. H. Tran, T. D. Nguyen, W. Bao, A. R. Balef, B. B. Zhou, and A. Y. Zomaya, "DONE: Distributed approximate Newton-type method for federated edge learning," *IEEE Transactions on Parallel and Distributed Systems*, vol. 33, no. 11, pp. 2648–2660, 2022.
- [14] Z. Yao, A. Gholami, S. Shen, M. Mustafa, K. Keutzer, and M. Mahoney, "Adahessian: An adaptive second order optimizer for machine learning," in *proceedings of the AAAI conference on artificial intelligence*, vol. 35, no. 12, 2021, pp. 10665–10673.
- [15] M. F. Hutchinson, "A stochastic estimator of the trace of the influence matrix for laplacian smoothing splines," *Communications in Statistics-Simulation and Computation*, vol. 18, no. 3, pp. 1059–1076, 1989.
- [16] H. Liu, Z. Li, D. Hall, P. Liang, and T. Ma, "Sophia: A scalable stochastic second-order optimizer for language model pre-training," *arXiv preprint arXiv:2305.14342*, 2023.
- [17] S. Lee, A. K. Sahu, C. He, and S. Avestimehr, "Partial model averaging in federated learning: Performance guarantees and benefits," *Neurocomputing*, vol. 556, p. 126647, 2023.
- [18] J. Martens, I. Sutskever, and K. Swersky, "Estimating the hessian by back-propagating curvature," *arXiv preprint arXiv:1206.6464*, 2012.
- [19] A. Ghalkha, C. Ben Issaid, A. Elgabli, and M. Bennis, "DIN: A decentralized inexact Newton algorithm for consensus optimization," in *ICC 2023 - IEEE International Conference on Communications*, 2023, pp. 4391–4396.