

FedSL: Federated Split Learning for Collaborative Healthcare Analytics on Resource-Constrained Wearable IoMT Devices

Wanli Ni, Huiqing Ao, Hui Tian, Yonina C. Eldar, *Fellow, IEEE*, and Dusit Niyato, *Fellow, IEEE*

Abstract—Many wearable Internet of Medical Things (IoMT) devices have limited computing power and small storage space. Additionally, the healthcare data sensed by a single IoMT device is not enough to train a sophisticated deep learning model. To address these challenges, we propose a federated split learning (FedSL) framework that allows for collaborative healthcare analytics on multiple IoMT devices with limited resources. Compared to centralized learning, FedSL can protect user privacy by not sending raw data over wireless networks. Furthermore, FedSL offers more flexibility than other federated learning methods. It enables even low-end IoMT devices to participate in model training and result inference. Experimental results show that our FedSL performs well on medical imaging tasks with different data distributions.

Index Terms—Federated split learning, healthcare analytics, Internet of Medical Things, user privacy, wearable devices.

I. INTRODUCTION

The Internet of Medical Things (IoMT) enables real-time body data collection in a cost-effective way, which is changing existing healthcare industries [1]. Moreover, with the success of deep learning (DL) in medical diagnosis tasks, online healthcare platforms can now train powerful models using data from massive IoMT devices [2], [3]. This enables patients to get telemedicine services quickly without visiting a clinic, and doctors can make more reliable and faster decisions [4]. But how to use wearable IoMT devices for collaborative healthcare analytics is still a challenge, especially with limited resources.

Due to the highly sensitive nature of medical information, sending IoMT data over wireless networks may pose a significant privacy concern. Recently, federated learning (FL) has emerged as a potential solution to overcome this issue [5], [6]. However, FL requires all devices to train the full DL model, which can be resource-intensive, especially for models with millions of parameters [5], [6]. This is impractical for many small-size and low-cost wearable IoMT devices, such as patch sensors and smart watches [1]. To this end, the split learning paradigm can be adopted to divide the complete DL model into a server-side model and a device-side model [2]. By doing

The work of Hui Tian was supported by the Beijing Municipal Natural Science Foundation under Grant L232052. (*Corresponding author: Wanli Ni.*)

Wanli Ni is with the Department of Electronic Engineering, Tsinghua University, and also with the Beijing National Research Center for Information Science and Technology, Beijing 100084, China (e-mail: niwanli@tsinghua.edu.cn).

Huiqing Ao and Hui Tian are with the State Key Laboratory of Networking and Switching Technology, Beijing University of Posts and Telecommunications, Beijing 100876, China (e-mail: {hqao, tianhui}@bupt.edu.cn).

Yonina C. Eldar is with the Faculty of Mathematics and Computer Science, Weizmann Institute of Science, Rehovot 7610001, Israel (e-mail: yonina.eldar@weizmann.ac.il).

Dusit Niyato is with the School of Computer Science and Engineering, Nanyang Technological University, Singapore (e-mail: dniyato@ntu.edu.sg).

The simulation code and extended analysis for this paper are available at <https://github.com/niwanli/FedSL-IoMT>.

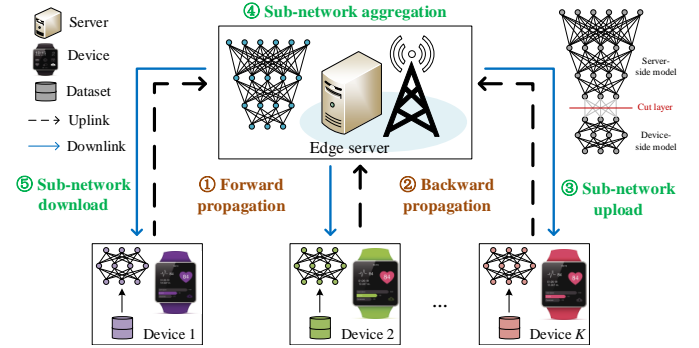


Fig. 1. An illustration of the proposed FedSL framework for IoMT devices.

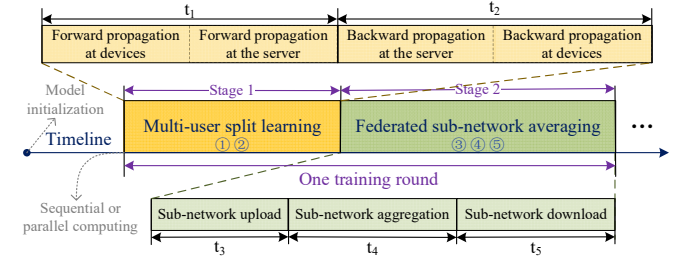


Fig. 2. One training round of FedSL in the IoMT-based healthcare system.

so, IoMT devices only need to train a small portion of the DL model, effectively reducing their computational burden [7]. As a federated variant of split learning, in this paper, we propose a federated split learning (FedSL) framework to alleviate local hardware requirements and support semi-distributed learning while partially preserving user privacy.

II. FRAMEWORK DESIGN

As illustrated in Fig. 1, we propose a FedSL framework for resource-limited wearable IoMT devices to train a shared DL model collaboratively using the medical healthcare data collected. The set of devices is denoted as $\mathcal{K} = \{1, 2, \dots, K\}$. The layers of the DL model are denoted as $\mathcal{L} = \{L_1, L_2, \dots, L_N\}$. According to the depth of neural networks, the full DL model is divided into two sub-networks [2]. Specifically, these high-level layers $\mathcal{L}_S = \{L_{C+1}, L_{C+2}, \dots, L_N\}$ are kept at the server, while the low-level layers $\mathcal{L}_D = \{L_1, L_2, \dots, L_C\}$ are stored on IoMT devices, where L_C is the cut layer [3]. FedSL's model splitting paradigm offers several advantages: 1) Similar to FL, FedSL can fully utilize the distributed data collected by wearable IoMT devices while preserving user privacy. This is achieved by transferring only the intermediate features extracted by the device-side model to the server, enabling collaborative model training without disclosing sensitive raw data. 2) Unlike FL, which requires clients with sufficient computing and storage resources, FedSL enables computing-limited and storage-constrained IoMT devices to participate in the model




Fig. 3. Normal chest X-Ray image (left) versus viral pneumonia image (right).

training process. This is because only shallow neural networks are computed on these devices, reducing computational and storage requirements. Compared to the scheme in [7], FedSL only needs one edge server to coordinate model training, and thus the communication flow is different. Specifically, as shown in Fig. 2, each training round of FedSL has two stages.

- **Stage 1 (Multi-user split learning):** At the beginning, all devices feed sub-networks with mini-batch data samples in sequence or parallel. Next, the smashed data obtained at the cut layer and the label are sent to the edge server for completing the rest of the forward propagation [7]. Then, with the loss values obtained at the output layer, the gradients are computed for different devices and back propagated from the server to devices.
- **Stage 2 (Federated sub-network averaging):** In the sequel, to share knowledge and accelerate convergence, all devices upload their updated sub-network parameters to the server. After aggregation, the averaged sub-network parameters are sent back to all devices for the next training round [5], [6]. The learning process terminates when the loss value stabilizes or the maximum number of training rounds is reached.

III. EXPERIMENTAL RESULTS

We conduct simulations in an IoMT-based healthcare monitoring system with $K = 5$ devices. Specifically, an 18-layer ResNet is employed to classify Chest X-Ray (two classes as shown in Fig. 3) and Optical Coherence Tomography (OCT, four classes) images [4] with independent identically distributed (IID) and non-IID settings. The cut layer is set as $L_c = 3$, the learning rate is 10^{-4} , and the maximum number of training rounds is 400. For comparison, we consider four schemes: i) Centralized learning (CL), ii) FL, iii) Sequential FedSL, and iv) Parallel FedSL.

Fig. 4 shows the learning performance of training ResNet-18 on the Chest X-Ray dataset with IID and non-IID settings. From this figure, we have the following observations: 1) In the first half of the training process, the sequential FedSL scheme is able to achieve similar performance to CL, while the parallel FedSL scheme is close to the FL baseline. This is because in the sequential FedSL scheme, the server only needs to train one server-side model for all devices, while in the parallel FedSL scheme, the server needs to compute different sub-networks for different devices in parallel. 2) When the number of training rounds is greater than 50, the performance gap between any two schemes is very small. This is due to the fact that the DL model is well trained in the second half of all schemes, and thereby all schemes achieve almost similar prediction accuracy. 3) Compared to the IID setting, these results in the non-IID case fluctuate more significantly and the




Fig. 4. Accuracy on the chest X-Ray dataset with IID and non-IID settings.




Fig. 5. Prediction accuracy on the OCT dataset with IID and non-IID settings.

sequential and parallel FedSL schemes slightly outperform the traditional CL and FL benchmarks, respectively. This shows that our FedSL has good adaptability to the non-IID datasets. Similar observations can be seen in Fig. 5, due to the page limit, we omit the analysis here for brevity. Please go to our site at <https://github.com/niwanli/FedSL-IoMT> for more details. Overall, the findings above reveal the effectiveness and superiority of the proposed FedSL in healthcare systems.

IV. CONCLUSION

In this paper, we proposed the FedSL framework to implement medical data analysis on multiple wearable IoMT devices. This framework is friendly to resource-constrained, low-cost terminals in comparison to traditional CL and FL schemes, especially when the size of the trained DL model is large. Furthermore, since only intermediate features are shared between servers and devices, our FedSL can protect user privacy partially without compromising data accessibility, thus making it suitable for privacy-sensitive cases. Experiments showed that our FedSL achieves the similar performance as FL, and the gap between FedSL and CL is small.

REFERENCES

- [1] A. Gatouillat *et al.*, “Internet of Medical Things: A review of recent contributions dealing with cyber-physical systems in medicine,” *IEEE Internet of Things J.*, vol. 5, no. 5, pp. 3810–3822, Oct. 2018.
- [2] P. Vepakomma *et al.*, “Split learning for health: Distributed deep learning without sharing raw patient data,” Dec. 2018. [Online]. Available: <https://arxiv.org/pdf/1812.00564.pdf>
- [3] M. Zhang *et al.*, “SplitAVG: A heterogeneity-aware federated deep learning method for medical imaging,” *IEEE J. Biomed. Health Informatics*, vol. 26, no. 9, pp. 4635–4644, Sept. 2022.
- [4] D. S. Kermany *et al.*, “Identifying medical diagnoses and treatable diseases by image-based deep learning,” *Cell*, vol. 172, no. 5, pp. 1122–1131.e9, Feb. 2018.
- [5] T. Gafni *et al.*, “Federated learning: A signal processing perspective,” *IEEE Signal Process. Mag.*, vol. 39, no. 3, pp. 14–41, May 2022.
- [6] L. You *et al.*, “A triple-step asynchronous federated learning mechanism for client activation, interaction optimization, and aggregation enhancement,” *IEEE Internet of Things J.*, vol. 9, no. 23, pp. 24 199–24 211, Dec. 2022.
- [7] C. Thapa *et al.*, “SplitFed: When federated learning meets split learning,” in *Proc. AAAI*, Feb. 2022, pp. 8485–8493.