

Balancing Uneven Knowledge of Hospital Nodes for ICU Patients Diagnosis through Federated Learning

Claudia Di Napoli¹, Giovanni Paragliola^{1,*}, Patrizia Ribino¹ and Luca Serino¹

¹National Research Council, Institute for High-Performance Computing and Networking (ICAR), Italy

Abstract

The Covid pandemic highlighted the urgent need for collaborations in the healthcare sector to empower clinical and scientific communities in responding to global challenges. In this context, the ICU4Covid project joined research institutions, medical centers, and hospitals all around Europe in a European Telemedicine Network, allowing for sharing of capabilities, knowledge, and expertise distributed in such a network. Nevertheless, healthcare data sharing has ethical, regulatory, and legal complexities imposing restrictions on access and use. In addition, data and knowledge are very often unevenly distributed at the different nodes of the network depending on their geographical location and dimension. To address these issues, a federated learning architecture is proposed to allow for distributed machine learning within the cross-institutional healthcare system without moving data outside its original location. The approach has been applied for the early prediction of high-risk hypertension patients. The experimentation carried out shows that the knowledge of single nodes is spread within the federation, improving the ability of each of them to perform predictions also on not previously treated cases. The performance evaluation of the computed predictions in terms of accuracy and precision is over 0.91 confirming the encouraging results of the proposed FL approach.

Keywords

Federating Learning, Predictive Models for Healthcare, Telemedicine Network

1. Introduction

SARS-CoV-2 pandemic highlights the need for improving cooperation and knowledge sharing to prevent disease spread and ensure patient care quality. The pandemic showed that the uneven distribution of capacities and resources between healthcare organizations situated in small centers and those in urban areas makes it difficult to provide the same quality of healthcare services. To address these challenges, a network of research institutions, medical centers, and hospitals all around Europe join under the umbrella of the ICU4Covid project.

The ICU4Covid project [1] aims to create the sense of being part of the European telemedicine network composed of a set of independent Cyber-Physical Systems for Telemedicine and Intensive Care (CPS4TIC). It also aims to access the network's capabilities, knowledge, and expertise. Moreover, during the pandemic, the lack of large-scale healthcare organization intelligence put in more evidence the need for extensive and varied data sets for training ML algorithms for clinical purposes. In

this context, the ICU4Covid project provides valuable ground for learning from real-world health data that has proven to be effective in multiple healthcare applications, resulting in improved quality of care [2] [3], predicting disease risk factors [4][5], and analyzing genomic data for personalized medicine [6].

However, a healthcare ecosystem should address the problem that accessing or sharing health data outside the host institution is restricted by regulatory policies mandated by EU General Data Protection Regulation (GDPR) [7]. Thus, traditional or centralized machine learning algorithms, which require aggregating such distributed data into a central repository for the purpose of training a model, cannot be exploitable. Leveraging such data while complying with data protection policies requires rethinking data analytics methods for healthcare applications. In order to guarantee the sharing of knowledge between each node of the telemedicine network, we integrated a Federated Learning (FL) architecture in each node of the CPS4TIC system. The FL architecture enables the individual nodes of the network to act as local learners and send local model parameters to a central server instead of training data, so individual nodes independently train and collaboratively learn models without sharing local datasets. The central server aggregates the local models, defining a single global model, which is sent back to the clients to proceed with the FL process until all rounds are completed.

The proposed approach allows for balancing data intelligence so small and medium-sized healthcare organizations can benefit from collective intelligence without

Ital-IA 2023: 3rd National Conference on Artificial Intelligence, organized by CINI, May 29–31, 2023, Pisa, Italy

*Corresponding author.

✉ claudia.dinapoli@icar.cnr.it (C. Di Napoli);
giovanni.paragliola@icar.cnr.it (G. Paragliola);
patrizia.ribino@icar.cnr.it (P. Ribino); luca.serino@icar.cnr.it
(L. Serino)

📞 0000-0002-8626-5805 (C. Di Napoli); 0000-0003-3580-9232
(G. Paragliola); 0000-0003-3266-9617 (P. Ribino);
0000-0003-0077-1799 (L. Serino)

© 2023 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).
CEUR Workshop Proceedings (CEUR-WS.org)



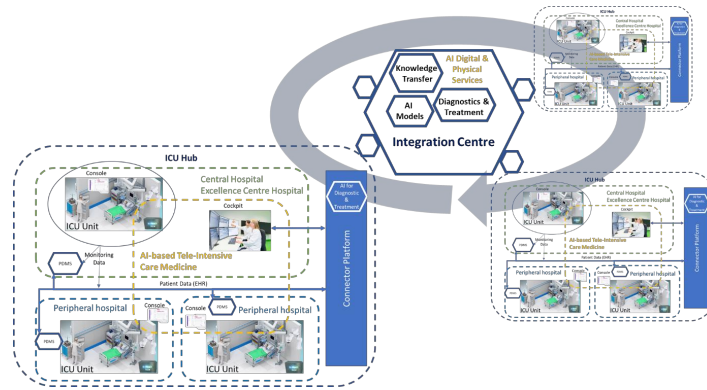


Figure 1: ICU4Covid Telemedicine network

requiring large data sets. We show how the knowledge owned by a single organization is spread among all the members of the federation by improving the reliability of the local model as a prediction test. A quantitative and qualitative estimation of improvement brought by the federated process over the local nodes is provided, reporting an enhancement of the performance up to 38%. Finally, a comparison between federated learning and the centralized approach shows that the federated approach prevents data privacy and security issues at a slight performance loss.

The rest of the paper is organized as follows. Section 2 introduces the ICU4COVID European Project. Section 3 shows an application scenario of federated learning for enhancing healthcare knowledge. Section 4 presents the performance evaluation. Finally, in Section 5, conclusions and future works are presented.

2. Overview of ICU4COVID European Project

The pandemic caused by the SARS-CoV-2 and subsequently by the COVID-19 virus showed that the uneven distribution of capacities and resources of Intensive Care Units (ICUs) located in rural and urban areas remains a big challenge. Hence, real-time information sharing and cooperation between hospitals, healthcare workers, and the public are highly significant to containing Covid-19 spreading and ensuring high-quality healthcare services. The Cyber-Physical System for Telemedicine and Intensive Care (CPS4TIC) aims at expanding Information Technology-based operations and information sharing from the central ICU Hubs to peripheral or rural hospitals while substantially minimizing the infection risk for healthcare staff (see a conceptual view in Fig.1). The CPS4TIC framework comprises a telemed-

ical cockpit, a telemedicine console installed in every peripheral hospital, a connector platform, and smart bedside hubs. The ICU4Covid project is aimed to deploy the CPS4TIC in many hospitals across Europe as a global network. ICU4Covid project advances the CPS4TIC to large-scale experimentation and deployment with full-scale participation of healthcare staff, hospitals, and end-users. Adapting the aforementioned innovative technology enables contemporary Intensive Care Units transformation into a structure that operates as one ICU Hub consisting of one centralized hospital connected to its peripheral hospitals in a geographical area. Each ICU Hub is composed of a central ICU and interconnected peripheral hospitals employing telemedicine and telemonitoring techniques that assist healthcare staff in patient screening, diagnosing, and treatment. Each ICU Hub is equipped with state-of-the-art technology, such as a 5G module, radar sensors, and AI chips, and controlled by a control station called Integration Center. The ICU4Covid architecture allows for the coexistence of both already established and new ICUs as one ICU node. In fact, the system is independent of the hospital's infrastructure enabling highly encrypted telemedicine and digitalization of ICUs with collective technological efforts.

3. Federated Learning for enhancing healthcare knowledge

To allow for sharing knowledge between each ICU Hub without incurring data privacy and security risks, we propose to equip each node of the CPS4TIC system with a federated learning (FL) architecture. The working scenario consists of three participants respectively named *Hospital 1*, *Hospital 2*, and *Hospital 3* with the same data

structure of the CPS4TIC hosted in each hospital, as illustrated in Fig. 2. The data of a participant is private and owned by the hospital, and it is used for local training to learn a local model. Each participant uses a local hybrid network, where hybrid means that the network is composed of different deep learning networks. Each local model updates from each participant are sent to an aggregator server that combines them into a global consensus model. This global model is then returned to each participant for further local training. The participants connect to the aggregator server through remote procedure calls via a transport layer security network connection. Sensitive information such as model, optimizer weights and aggregated metrics move between the participant and the aggregator server over this encrypted channel. For the sake of experimentation, these three participants are emulated on local clusters.

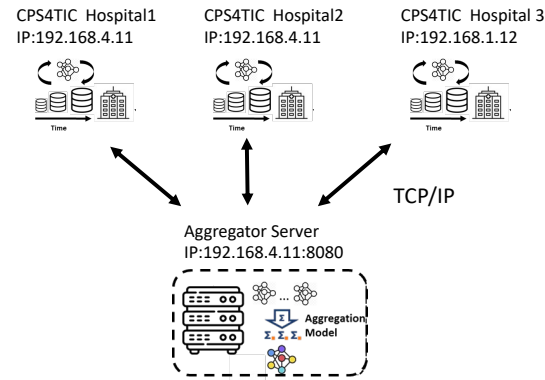


Figure 2: Federated CPS4TIC nodes

The proposed approach supports the decision-making process of a network of federated hospitals, as those provided by the ICU4Covid project. In order to validate the approach, it is applied to a use-case scenario for predicting high-risk hypertensive patients.

An ECG sample of a patient is taken from the dataset belonging to CPS4TIC of client 1 (CPS4TIC-1) stored only on CPS4TIC-1, named $Patient_x$. Such a sample is labelled as *high-risk* class. Therefore, CPS4TIC-2 and CPS4TIC-3 do not learn anything specifically from that patient. Fig. 3 shows the progression of the federated training process. In the beginning, only CPS4TIC-1 is able to correctly classify the validation sample, while CPS4TIC-2 and CPS4TIC-3 can not. At this stage, also the aggregated model fails to correctly classify the sample since experimental results show that more training rounds are necessary to successfully merge the knowledge from the local models. In the middle of the federated learning process, CPS4TIC-1 still correctly classifies the validation sample, as well as the aggregated model be-



Figure 3: Logical overview of the federated process validation

cause the knowledge of the CPS4TIC-1 model has been included in the aggregated model according to the federated learning process.

At the end of the federated learning process, all models can correctly classify the validation sample since the aggregation server shares the updated parameters round-by-round with the clients involved in training, so collecting the knowledge of all nodes.

Table 1 reports the qualitative accuracy trend during the training process showing how the accuracy changes in the different stages of the training.

Locally at CPS4TIC 1, the accuracy consistently achieves good values from the beginning of the training, while the other nodes and the aggregated model expose low accuracy values ranging from 52% to 61%. As the rounds progress, CPS4TIC 1 and the aggregator models improve their performance to 89% and 75%, respectively. Following the FL algorithm, the aggregator model's knowledge is spread to CPS4TIC 2 and CPS4TIC 3, and consequently, at the end of the training process, their performance improved, achieving 84% and 81%, respectively.

Without the federated learning process, Hospital 2 or Hospital 3 would have classified $Patient_x$ as a low-risk patient and consequently, no healthcare protocol would have been adopted for that patient. Conversely, Hospital 2 and Hospital 3 can correctly classify $Patient_x$ allowing them to adopt appropriate healthcare since their local models are updated with the knowledge of the aggregated model.

The improvement of the local models due to the federated process is shown in Table ?? reporting an increase from the beginning to the end of the training equals 38% and 35%, respectively.

	Local Model CPS4TIC 1	Local Model CPS4TIC 2	Local Model CPS4TIC 3	Aggregator Model
Start Training	0.85	0.52	0.52	0.61
Middle Training	0.89	0.62	0.61	0.75
End Training	0.92	0.84	0.81	0.91

Table 1
Accuracy trend during the training process stages

4. Performance evaluation

This section provides a performance evaluation of the federated learning approach compared with the classical centralized one. The evaluation was conducted using the SHAREE [27] database and considering three local nodes for the federation. The SHAREE database contains 169 electrocardiographic (ECG) records of hypertensive patients monitored with an epicardial holter for 24 hours with an attempt to record major cardiovascular events. Patients who experienced dangerous events were marked as *high-risk*, while the rest were marked as *low-risk*. In such an experimental evaluation, the dataset is used to train machine learning models to identify subjects at a higher risk of developing fatal cardiovascular events. The learning process uses multivariate time series (MTS) data, whose raw signal comes from three electrodes placed on the subject's chest during monitoring.

Five-minute segments of input data are randomly selected as samples. Hence, the training set contains more than 14,000 samples evenly distributed between the two classes, *high-risk* and *low-risk*. Training set samples were equally distributed among the three local nodes of the federation. The test set, instead, is defined using the hold-out approach. Thus, 700 samples are used on each client node, with two-thirds marked as *low-risk* and one-third as *high-risk*.

Table 2 shows the performance achieved in terms of accuracy by each local model on the test set at the end of the learning process. As it is possible to note, the aggregated model achieves slightly higher performance than the local models. Indeed, the accuracy of the aggregated model is equal to 90% higher than each local model, 87%, 88%, and 88%, respectively.

	Local Model1	Local Model2	Local Model3	Aggregated Model
Acc	0.87	0.88	0.88	0.90

Table 2
Comparison between local models and the aggregated model

Table 3 compares the federated and classical centralized approaches in terms of accuracy and precision. As it is possible to note, the *Federated approach* achieves an accuracy of 90% and a Precision of 91%, respectively.

Approach	Accuracy	Precision
Federated Scenario	0.90+-0.0019	0.91+-0.0059
Centralised Scenario	0.98+-0.005	0.98+-0.002

Table 3
Comparison between federated and centralized approach

Whilst, the *Centralized approach* has achieved the best performance, with Accuracy and Precision values of 98%.

Despite the better performance of the centralized approach, the performance results achieved by the federated model can be considered satisfactory in terms of Accuracy and Precision, considering the advantages coming from the adoption of a federated approach in the healthcare domain. The centralized mode will require moving all data from its stored nodes to the node performing the learning process. Thus, data security and privacy are compromised by this action. Fig. 4 clearly shows that in the centralized approach, 100% of data are moved across the nodes to be collected in a unique node. On the contrary, the federated approach prevents privacy risks since no data are moved; only the federated model parameters are transferred.




Figure 4: % of data moved during the model training

For the truth's sake, we must recall that the federated approach can introduce communication costs issues while sharing the neural network parameters, issues evaluated and addressed in [25, 26].

Our results highlight a trade-off between performance and security, reported in Fig. 5, which visually describes the relationship between these two aspects for a generic predictive model. The figure defines a qualitative visual




Figure 5: Graphical View of Trade-off between Performance and Security

space delimited by the projection of privacy and performance. It does not define quantitative values since the purpose is only to provide a simple and immediate tool for comparison between centralized and federated models. The upper right corner explains this trade-off in terms of traditional centralized approach-based predictive models characterized by high performance and high data security risk. On the other hand, from a federated point of view, this approach achieves good levels of performance, not like those centralized, but the risk of data protection is low, close to zero. Our research reveals a target in the lower right corner where performance is comparable to centralized and privacy risk is low.

5. Conclusion

The paper presents a federated learning approach to support the diagnosis process of hypertensive patients in a European telemedicine network. A network of federated medical institutions demonstrates how each organization's knowledge is disseminated to all association members by enhancing the quality of results and the reliability of local models as predictive tests. Finally, conventional metrics quantify the performance of the proposed approach compared to the centralized one. Results confirm that the FL approach can significantly support the decision-making process of intensive care patients in distributed networks of federated medical organizations.

Moreover, experimental results highlighted new issues for future research, such as the evaluation of the impact of new clients and larger datasets, as well as the definition of the trustworthiness of a new client involved in the training process.

Acknowledgments

This work is based on results obtained within the Cyber-Physical Intensive Care Medical System for Covid-19 (ICU4Covid) European Project (Grant Agreement number: 101016000).

References

- [1] ICU4Covid- Cyber-Physical Intensive Care Medical System for Covid-19 [Online]. Available: <https://www.icu4covid.eu/>. [Accessed: 23-Jan-2023].
- [2] Johnson, A. E., Ghassemi, M. M., Nemati, S., Niehaus, K. E., Clifton, D. A., & Clifford, G. D. (2016). Machine learning and decision support in critical care. *Proceedings of the IEEE*, 104(2), 444-466.
- [3] Chakrabarty, A., Zavitsanou, S., Sowrirajan, T., Doyle III, F. J., & Dassau, E. (2019). Getting IoT-ready: The face of next-generation artificial pancreas systems. In *The Artificial Pancreas* (pp. 29-57). Academic Press.
- [4] Poplin, R., Varadarajan, A. V., Blumer, K., Liu, Y., McConnell, M. V., Corrado, G. S., ... & Webster, D. R. (2018). Prediction of cardiovascular risk factors from retinal fundus photographs via deep learning. *Nature Biomedical Engineering*, 2(3), 158-164.
- [5] Zheng, T., Xie, W., Xu, L., He, X., Zhang, Y., You, M., Yang, G., & Chen, Y. (2017). A machine learning-based framework to identify type 2 diabetes through electronic health records. *International journal of medical informatics*, 97, 120-127.
- [6] Libbrecht, M. W., & Noble, W. S. (2015). Machine learning applications in genetics and genomics. *Nature Reviews Genetics*, 16(6), 321-332.
- [7] Voigt, P., & Von dem Bussche, A. (2017). *The EU general data protection regulation (GDPR). A Practical Guide*, 1st Ed., Cham: Springer International Publishing, 10(3152676), 10-5555.
- [8] Konečný J, McMahan HB, Ramage D, Richtárik P. Federated optimization: Distributed machine learning for on-device intelligence. arXiv preprint arXiv:1610.02527. 2016 Oct 8.
- [9] Yang Q, Liu Y, Chen T, Tong Y. Federated machine learning: Concept and applications. *ACM Transactions on Intelligent Systems and Technology (TIST)*. 2019 Jan 28;10(2):1-9.
- [10] Sellamna A, Boukhamla AZ, Benkaddour MK. Federated Learning paradigm in E-health systems: An overview. In *2022 4th International Conference on Pattern Analysis and Intelligent Systems (PAIS) 2022 Oct 12* (pp. 1-5). IEEE.
- [11] Chen H, Li H, Xu G, Zhang Y, Luo X. Achieving privacy-preserving federated learning with irrelevant updates over e-health applications. *InICC 2020-*

- 2020 IEEE International Conference on Communications (ICC) 2020 Jun 7 (pp. 1-6). IEEE.
- [12] Fourati LC, Samiha AY. Federated learning toward data preprocessing: COVID-19 context. In 2021 IEEE International Conference on Communications Workshops (ICC Workshops) 2021 Jun 14 (pp. 1-6). IEEE.
- [13] Anavangot V, Lukose J. Calibration and Classifier Design for IoT Healthcare Applications using Federated Learning. In 2022 International Conference on Computing, Communication, Security and Intelligent Systems (IC3SIS) 2022 Jun 23 (pp. 1-6). IEEE.
- [14] Han B, Jhaveri R, Wang H, Qiao D, Du J. Application of robust zero-watermarking scheme based on federated learning for securing the healthcare data. *IEEE Journal of Biomedical and Health Informatics*. 2021 Nov 12.
- [15] Lai J, Song X, Wang R, Li X. Edge intelligent collaborative privacy protection solution for smart medical. *Cyber Security and Applications*. 2023 Dec 1;1:100010.
- [16] Wang H, Zhang X, Xia Y, Wu X. An intelligent blockchain-based access control framework with federated learning for genome-wide association studies. *Computer Standards & Interfaces*. 2023 Mar 1;84:103694.
- [17] Aich S, Sinai NK, Kumar S, Ali M, Choi YR, Joo MI, Kim HC. Protecting personal healthcare record using blockchain & federated learning technologies. In 2022 24th International Conference on Advanced Communication Technology (ICACT) 2022 Feb 13 (pp. 109-112). IEEE.
- [18] Florescu LM, Streba CT, Șerbănescu MS, Mămuleanu M, Florescu DN, Teică RV, Nica RE, Gheonea IA. Federated Learning Approach with Pre-Trained Deep Learning Models for COVID-19 Detection from Unsegmented CT images. *Life*. 2022 Jun 26;12(7):958.
- [19] Durga R, Poovammal E. FLED-Block: Federated Learning Ensembled Deep Learning Blockchain Model for COVID-19 Prediction. *Frontiers in Public Health*. 2022;10.
- [20] Linardos A, Kushibar K, Walsh S, Gkontra P, Lekadir K. Federated learning for multi-center imaging diagnostics: a simulation study in cardiovascular disease. *Scientific Reports*. 2022 Mar 3;12(1):1-2.
- [21] Wu C, Wu F, Lyu L, Huang Y, Xie X. Communication-efficient federated learning via knowledge distillation. *Nature communications*. 2022 Apr 19;13(1):1-8.
- [22] Nasser N, Fadlullah ZM, Fouda MM, Ali A, Imran M. A lightweight federated learning based privacy preserving B5G pandemic response network using unmanned aerial vehicles: A proof-of-concept. *Computer Networks*. 2022 Mar 14;205:108672.
- [23] Nguyen TV, Dakka MA, Diakiw SM, VerMilyea MD, Perugini M, Hall JM, Perugini D. A novel decentralized federated learning approach to train on globally distributed, poor quality, and protected private medical data. *Scientific Reports*. 2022 May 25;12(1):1-2.
- [24] Das P, Singh M, Roy DG. Blockchain-Based COVID-19 Detection Framework Using Federated Deep Learning. In International Conference on Network Security and Blockchain Technology 2022 (pp. 369-378). Springer, Singapore.
- [25] Giovanni Paragliola. Evaluation of the trade-off between performance and communication costs in federated learning scenario. *Future Generation Computer Systems*. 2022. DOI: <https://doi.org/10.1016/j.future.2022.06.006>
- [26] Giovanni Paragliola, Antonio Coronato. Definition of a novel federated learning approach to reduce communication costs. *Expert Systems with Applications*. 2022. DOI: <https://doi.org/10.1016/j.eswa.2021.116109>
- [27] P. Melillo, R. Izzo, A. Orrico, P. Scala, M. Attanasio, M. Mirra, N. De Luca, L. Pecchia "Automatic prediction of cardiovascular and cerebrovascular events using Heart Rate Variability analysis", *Plos One*, March 20, 2015. <http://journals.plos.org/plosone/article?id=10.1371/journal.pone.0118504>.