

# Attention on Personalized Clinical Decision Support System: Federated Learning Approach

Chu Myaet Thwal, Kyi Thar, Ye Lin Tun, Choong Seon Hong  
*Department of Computer Science and Engineering*  
*Kyung Hee University*  
 Yongin-si, 17104, Republic of Korea  
 {chumyaet, kyithar, yelintun, cshong}@khu.ac.kr

**Abstract**—Health management has become a primary problem as new kinds of diseases and complex symptoms are introduced to a rapidly growing modern society. Building a better and smarter healthcare infrastructure is one of the ultimate goals of a smart city. To the best of our knowledge, neural network models are already employed to assist healthcare professionals in achieving this goal. Typically, training a neural network requires a rich amount of data but heterogeneous and vulnerable properties of clinical data introduce a challenge for the traditional centralized network. Moreover, adding new inputs to a medical database requires re-training an existing model from scratch. To tackle these challenges, we proposed a deep learning-based clinical decision support system trained and managed under a federated learning paradigm. We focused on a novel strategy to guarantee the safety of patient privacy and overcome the risk of cyberattacks while enabling large-scale clinical data mining. As a result, we can leverage rich clinical data for training each local neural network without the need for exchanging the confidential data of patients. Moreover, we implemented the proposed scheme as a sequence-to-sequence model architecture integrating the attention mechanism. Thus, our objective is to provide a personalized clinical decision support system with evolvable characteristics that can deliver accurate solutions and assist healthcare professionals in medical diagnosing.

**Index Terms**—clinical decision support system, healthcare, artificial intelligence, sequence-to-sequence network, attention mechanism, federated learning

## I. INTRODUCTION

Remarkable advances in digital and Internet-of-Things (IoT) technologies bring great potential to transform a wide range of cities into smart cities. With the rise of the smart cities concept, automation strategy, based on a tremendous deployment of IoT devices and software solutions, helps to improve cities' services and makes more efficient use of their existing infrastructures and assets [1]. Smart cities take the advantages of information and communications technology (ICT) to enhance a better quality of life for their citizens by delivering smarter infrastructures in education, transportation, healthcare, economy, and environmental management [2]–[4].

This work was partially supported by Institute of Information & communications Technology Planning & Evaluation (IITP) grant funded by the Korea government(MSIT) (No.2019-0-01287, Evolvable Deep Learning Model Generation Platform for Edge Computing) and by Institute of Information & Communications Technology Planning & Evaluation(IITP) grant funded by the Korea government(MSIT)(2020-0-00364, Development of Neural Processing Unit and Application systems for enhancing AI based automobile communication technology). \*Dr CS Hong is the corresponding author.

Thus, the goal of every smart city is to keep a balance in terms of sustainability in economic, environmental, and social impact [5]. Despite the countless benefits of smart city innovations, many challenges need to be focused on, as rapid urbanization introduces new problems to an already complex environment of a highly-populated area. Among them, health management has become one of the primary issues as new kinds of infectious diseases (i.e., Ebola, plague, COVID-19, etc.) are introducing to a rapidly growing modern society [6]. To prevent such disease outbreaks leading to the global pandemic, the authorities keep investigating the medical science and provide better innovations to the health management sector [7].

In recent years, artificial intelligence (AI) has contributed numerous kinds of innovations to the smart city domain and gains great attention in many different and popular applications [8]. AI also plays as a key component in developing better healthcare infrastructures ranging from patient monitoring to clinical data mining. As many AI algorithms are capable of learning from data, various types of supervised machine learning models can provide AI-based solutions for medical diagnoses and deliver accurate results to healthcare professionals [9]. A sufficient amount of clinical data is an important prerequisite for accelerating and promoting healthcare services through AI technology. For training a better medical diagnosis system, clinical data such as electronic health records (EHRs) from different sources of the medical organization need to be collectively organized [10]. However, these data may involve not only the privacy of individual patients but also the experimental details of medical organizations. Due to the heterogeneous and vulnerable properties, dealing with limited clinical data becomes a huge obstacle for big data analysis to facilitate precision medicine while developing better healthcare infrastructures [11].

In this paper, we proposed a deep learning-based clinical decision support system that can provide healthcare professionals with accurate predictions of potential diseases according to patients' symptoms in their medical records. We integrate the aid of edge AI that allows the processing of AI algorithms locally on edge devices (i.e., healthcare professionals' personal computers or mobile phones) for reducing the cost of data communication to perform the real-time operations and avoiding the transmission of vulnerable clinical data for a privacy

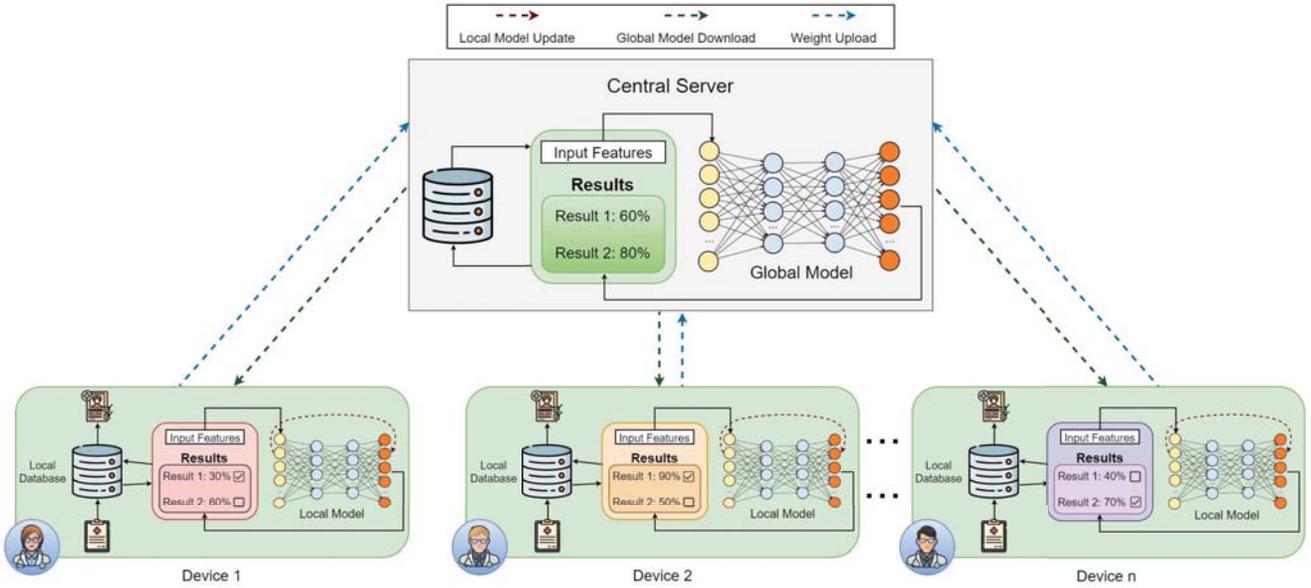


Fig. 1: System overview.

perspective [12]. We exploit the federated learning paradigm that can work with a certain degree of non-IID (Independent and Identically Distributed) data from heterogeneous sources by learning a shared global model at the central server while keeping all the sensitive clinical data on each local device [13]. As a result, local models can benefit from a global model with a predictive capacity of training on a rich source of clinical data while not compromising their data as well as preventing the cause of misdiagnosis.

Moreover, most of the existing medical diagnosis models need to be re-trained from scratch whenever new input features (i.e., symptoms) and target labels (i.e., diseases) are added in order to expand the prediction for new diseases [14], [15]. In our proposed system, we designed a sequence-to-sequence network for training local medical diagnosis models in which symptoms and diseases are treated as input and output sequences, respectively. For the models to be able to focus more on the relevant input sequences and predict accurate output sequences, we adopt the attention mechanism in designing the network architecture [16]. Thus, our system can easily adapt new input features and target labels and evolve to predict new kinds of diseases from heterogeneous sources without re-training the models from scratch [17].

Our objective is to provide an evolvable clinical decision support system that can assist healthcare professionals in the process of medical diagnosis while preserving the patients' privacy. Hence, patients can also be provided with higher quality, safer, and more efficient healthcare services. The main contributions of this paper can be listed as follows:

- We design the system architecture of an edge-based personalized clinical decision support system to jointly work with AI technology.
- We apply the attention mechanism in sequence-to-

sequence network of our medical diagnosis model for each edge device.

- We exploit the federated learning framework to train our personalized models collaboratively in a distributed manner on the datasets of non-IID symptom-disease network.
- We adopt a centralized global model to capture the shared knowledge of all personalized models by aggregating their gradients, regardless of their private information.
- We also design our models with evolvable characteristics to be able to adapt to new kinds of symptoms and diseases without re-training from scratch.

The rest of this paper is organized as follows: Section II explains the system architecture and methodologies of the proposed clinical decision support system. Section III presents the experiment settings for training the models. Simulation results are discussed in Section IV. Finally, Section V presents our conclusions and future works.

## II. SYSTEM ARCHITECTURE

Let us consider a personalized healthcare architecture in which different medical diagnosis models jointly participate in learning the heterogeneous clinical data from the EHRs of regional healthcare professionals. The global model in the centralized server does not take part in the training process but captures the shared knowledge of all local models and monitors their training rounds. Each local model trains its own neural networks on its private training dataset of medical records. Each trained model is deployed onto the personal devices of individual healthcare professional to assist in classifying the diseases based on the related symptoms. Classification outputs from those models are analyzed by the healthcare professionals based on their expert knowledge to

evaluate the final result. The evaluated results of symptom-disease pairs are stored in the local database for the model calibration. The local model sends its model updates to the central server after each round of the training process. The server aggregates the collected local model updates and creates a new single global model. Then, the aggregated global model is distributed for replacing the local models and updating their performance. Fig. 1 shows an overview of our proposed clinical decision support system operating under the federated learning paradigm.

#### A. Attention on Edge AI-based Medical Diagnosis Model

We design the architecture of our proposed clinical decision support system as a deep recurrent neural network (Deep RNN). While constructing the local medical diagnosis models, we process the sequence-to-sequence modeling for symptom-disease mapping, where symptoms are input sequences as well as diseases are treated as output sequences. Each mapping is comprised of two RNN models for the encoder network and decoder network, respectively. We establish direct shortcuts between the input symptoms and the target diseases by applying the attention mechanism in the encoder-decoder network. The encoder network with hidden states  $h_s$  processes the input sequence  $x$  and each of the source hidden state  $\bar{h}_s$  is compared with the current target hidden state  $h_t$  by a score function in Bahdanau's additive style [18] which can be defined as

$$score(h_t, \bar{h}_s) = v_a^\top \tanh(W_1 h_t + W_2 \bar{h}_s), \quad (1)$$

where  $W_1$ ,  $W_2$  and  $v_a$  are the weight matrices. Then, the result is normalized and the attention weights  $\alpha_{ts}$  are produced by

$$\alpha_{ts} = \frac{\exp(score(h_t, \bar{h}_s))}{\sum_{s'=1}^S \exp(score(h_t, \bar{h}_{s'}))}. \quad (2)$$

Based on the calculated attention weights, a smaller dimensional representation of the input sequence, which is known as a context vector  $c_t$  is computed as the weighted average of the source states and can be defined by

$$c_t = \sum_s \alpha_{ts} \bar{h}_s. \quad (3)$$

The decoder network is initialized with the context vector and trained to generate an output sequence  $y$ . The context vector is concatenated with the current target hidden state and yields the attention vector  $a_t$  which can be derived by

$$a_t = f(c_t, h_t) = \tanh(W_c [c_t; h_t]). \quad (4)$$

The final attention vector is used to derive the softmax logit and loss. By processing our symptom-disease mapping problems as sequence-to-sequence models and applying the attention mechanism, our models do not need to be re-trained from scratch whenever new input features are added to the training data. The architecture of an attention network embedded in each local medical diagnosis model of our proposed system is shown in Fig. 2.

#### B. Collaboration of Federated Learning

For a neural network to work well, it is required to be trained on a rich data source. In the traditional training pipeline of a neural network, the training process is centralized and it becomes infeasible to communicate as the data to be collected grows larger. To alleviate the expensive communication costs together with security and privacy issues of the centralized training scheme, we collaborate the federated learning settings to our model architecture [19]. In the federated learning paradigm of our proposed system, the centralized server initially distributes a base global model  $\mathcal{M}_{global}$  to be employed in each personal device of the healthcare authority as a local model  $\mathcal{M}_k$  that is parameterized by the weight  $W_k$ . Each device stores the personalized symptom-disease knowledge dataset  $\mathcal{D}_k = \{X_k, Y_k\}$  in its local database, where  $X_k$  denotes the input features (i.e., symptoms) and  $Y_k$  denotes the related target labels (i.e., diseases), respectively. Each dataset  $\mathcal{D}_k$  is used to train the local medical diagnosis model  $\mathcal{M}_k : X_k \rightarrow Y_k$ . The training procedure for each local model can be defined as

$$W'_k = \arg \min_{W_k} L(\mathcal{M}_k, \mathcal{D}_k, W_k), \quad (5)$$

where  $W'_k$  denotes the updated weight after training and  $L(\cdot)$  is the general loss function of each local model argumented with model structure  $\mathcal{M}_k$ , training data  $\mathcal{D}_k$ , and the local weight parameter  $W_k$ . At the end of each training round, a local model sends its model updates  $W'_k$  back to the server. The central server aggregates all the weight updates collected from the local models and produces a new global model  $\mathcal{M}_{global}$  that is parameterized by the global weight  $W_g$ . The aggregation mechanism is known as the weighted average of local model parameters which can be defined as

$$W_g = \sum_{k=1}^m \frac{|\mathcal{D}_k|}{\sum_{k=1}^m |\mathcal{D}_{k+1}|}, \quad (6)$$

where  $m$  denotes the number of devices participating in the proposed system. The aggregated global model is distributed back to the local models after each round of global model updates. By training our RNN models under the federated learning settings, we can reduce the communication costs between the local models and the server. As well as the local clients are no longer required to reveal their sensitive data to the central server or between each other. The aforementioned procedure of the federated learning setting is listed in Algorithm 1.

### III. EXPERIMENT SETTINGS

In this section, we analyze the effectiveness of our proposed personalized clinical decision support system by demonstrating a proof-of-concept scenario.

#### A. Datasets

We use the manually preprocessed dataset [20] retrieved from a Disease-Symptom Knowledge Database [21] which is generated by an automated method based on the information

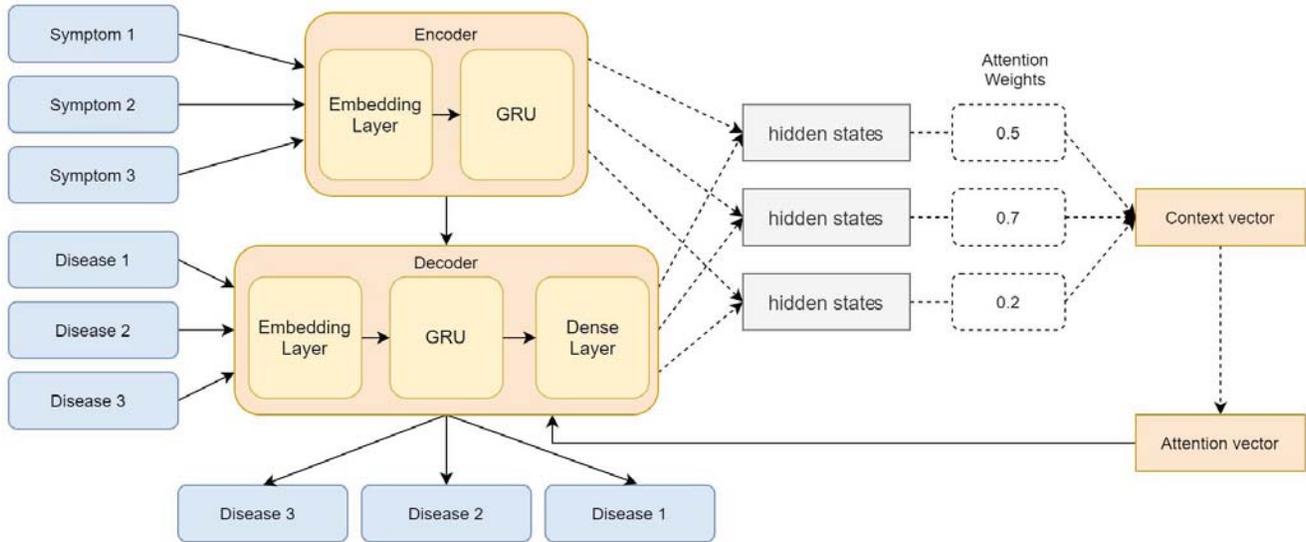


Fig. 2: Attention network architecture.

**Algorithm 1** Federated learning for global model update

- 1: **Inputs:** Local model  $M_k = \{M_1, M_2, \dots, M_K\}$  with parameter  $W_k$ , Global model  $M_{global}$  with parameter  $W_g$ , Input feature  $X_k$ , Target label  $Y_k$ , Local private datasets  $D_k = \{D_1, D_2, \dots, D_K\}$ , where each  $D_k = \{X_k, Y_k\}$
- 2: **Outputs:** Trained global model  $M_{global}$  with updated parameter  $W_g$
- 3: **Initialize:**  $W_g = random()$
- 4: **for** communication round  $t = 1, 2, \dots, T$  **do**
- 5:     **for** each client  $k = 1, 2, \dots, K$  in parallel **do**
- 6:          $W_k \leftarrow W_g$
- 7:         **for** local training step  $n = 1, 2, \dots, N$  **do**
- 8:             Train  $M_k : X_k \rightarrow Y_k$  with training data  $D_k$
- 9:             Update local model parameter:  

$$W'_k = \arg \min_{W_k} L(\mathcal{M}_k, \mathcal{D}_k, W_k)$$
- 10:         return updated local parameter  $W'_k$  to the central server
- 11:     **end for**
- 12:     **end for**
- 13:     Update global model parameter:  

$$W_g = \frac{1}{K} \sum_{k=1}^K W_k$$
- 14: **end for**
- 15: return  $W_g$

in textual discharge summaries of patients at New York Presbyterian Hospital in 2014. In the original database, the first column shows the types of diseases combined with UMLS codes obtained from MedLEE natural language processing system [22]. The second column shows the count of disease

occurrence. And, the third column shows the associated symptoms obtained by the statistical methods based on frequencies and co-occurrences. The original dataset is uncleaned and for the training processes to be more accurate, we use a manually preprocessed dataset that contains both training and testing sets of symptoms and diseases. The preprocessed dataset is cleaned and extensive, where the associations of symptoms are one hot encoded related to each disease. It contains 4,920 training samples having 41 unique types of diseases and 132 unique types of symptoms. The description of the preprocessed dataset is described in Table. I, where # denotes the number.

TABLE I: Statistics of dataset.

Data Description	
# of unique diseases	41
# of unique symptoms	132
# of clients	5
# of instances for client 1	1,000
# of instances for client 2	1,000
# of instances for client 3	1,000
# of instances for client 4	1,000
# of instances for client 5	920

We shuffle and split the training data into five subsets regarding local datasets for five clients. Four of the local datasets contain 1,000 instances and the remaining one contains 920 instances of diseases with related symptoms. We convert the datasets into text files in which symptoms and diseases are paired as input and output sentences. For the convenience of using with sequence-to-sequence networks, we add start and end tokens to each sentence and create a word index for mapping symptom to id and a reverse word index for mapping id to disease, respectively.

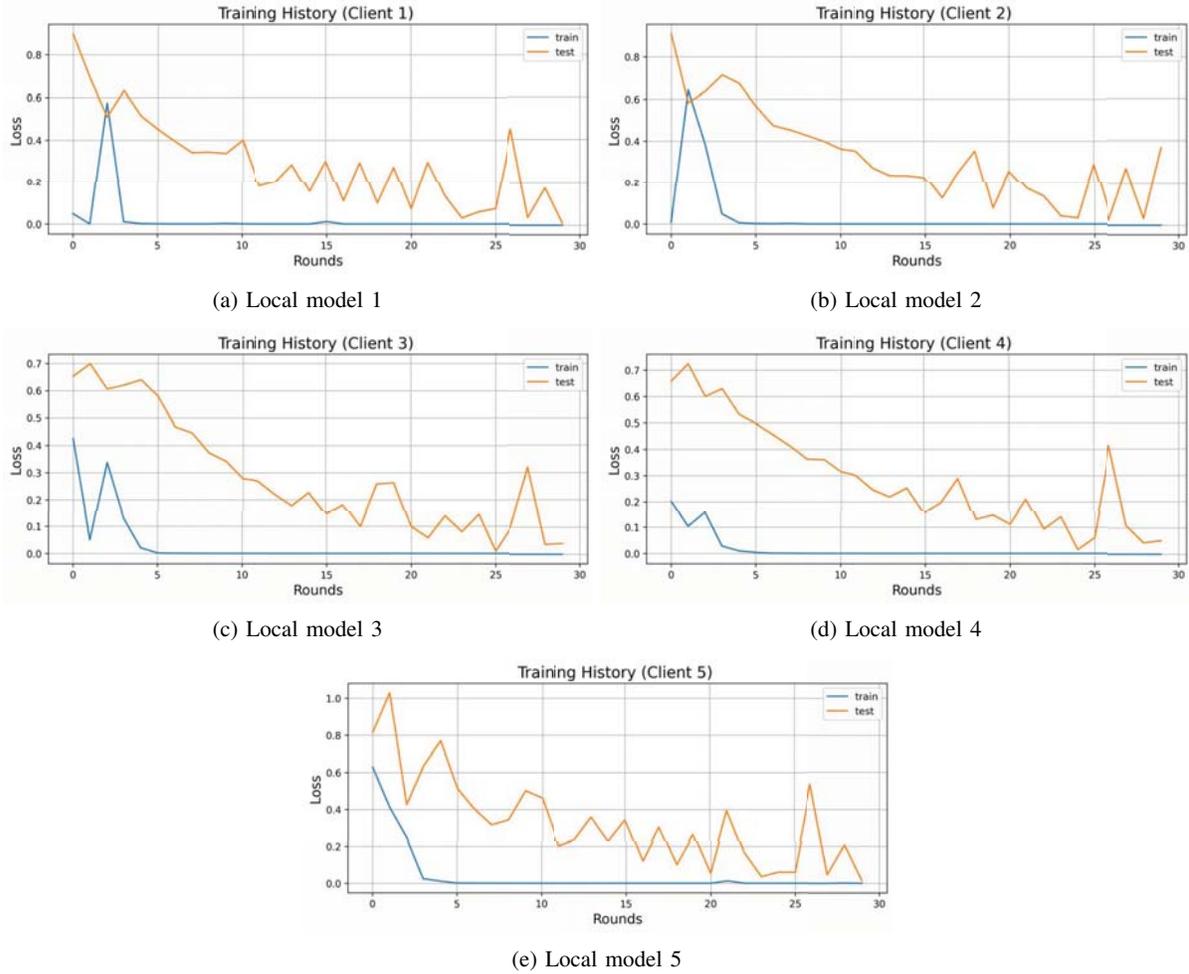


Fig. 3: Training history for centralized models.

## B. Local Models

We consider the personal devices of healthcare professionals as the clients for local models and construct one RNN model for each client of the local medical diagnosis system. For constructing the encoder-decoder network, we implement a gated recurrent unit (GRU) with 1,024 cells. Then, we apply the Bahdanau attention mechanism [23] to our neural networks.

## C. Training Settings

We train each local model on the dataset of the individual client in a distributed manner. For each client's data, we use 80% for training and 20% to validate the model. We randomize the initialization parameters of models and conduct five times of local training for each client per federated round. Then, we aggregate the weights of all local models by a weighted averaging mechanism and update the global model's weight. In the next round of federated training, all the local model weights are updated with the aggregated global model weight. We perform 30 times of federated training for a better accuracy of the disease prediction system.

## IV. SIMULATION RESULTS

In this section, we analyze the performance of our proposed clinical decision support system by executing the simulations on the Google Colab GPU backend. We consider a federated learning-based global healthcare architecture where regional healthcare professionals are participating as local clients with their own datasets. We develop the medical diagnosis models by using TensorFlow API [24]. We consider five clients participating as the local models in our healthcare architecture. Each local model is trained on the private dataset of an individual client. We illustrate the training and testing loss of each model as the benchmark centralized approach comparing to our proposed federated learning scheme.

The training history for the centralized models are shown in Fig. 3. As the individual local model is trained on its own local dataset and evaluates the results on the same dataset, all the local models converge after five rounds of the training process. However, the results of the testing processes are evaluated on the dataset of various clinical data. Thus, the testing losses of the local models fluctuate as the testing dataset includes

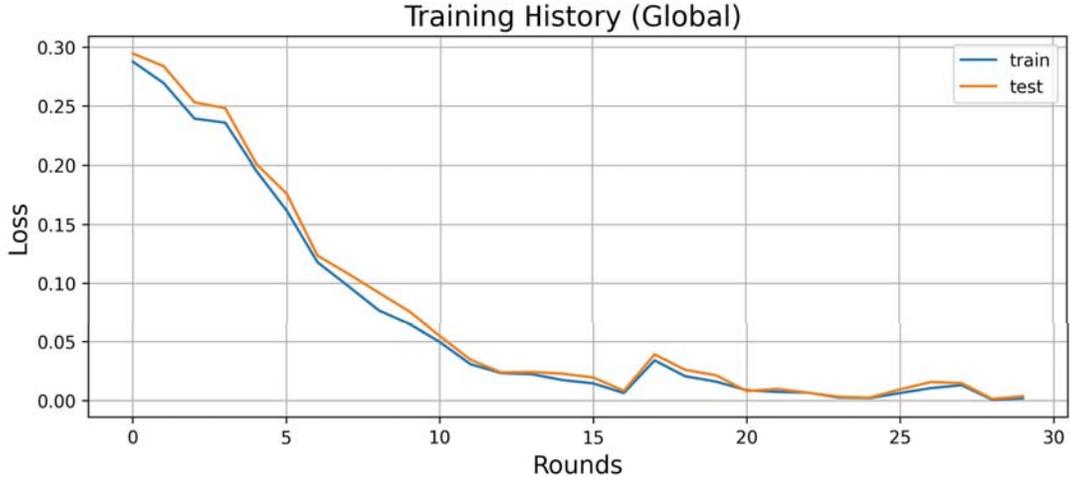


Fig. 4: Training history for global model.

unknown data for each local model. Fig. 3a shows the model performance of local client 1 where the training loss converges to the value of  $2.1083e-05$  and the testing loss varies from 0.0094 to 0.8976 throughout the whole simulation process. The model performance of local client 2 is shown in Fig. 3b where the training loss converges to the value of  $2.5834e-05$  and the testing loss varies from 0.0250 to 0.9088. Fig. 3c shows the model performance of local client 3 where the training loss converges to the value of  $2.0466e-05$  and the testing loss varies from 0.0089 to 0.6527. Fig. 3d and Fig. 3e show the model performances of local client 4 and local client 5, respectively. In Fig. 3d of local client 4, the training loss converges to the value of  $2.1468e-05$  and the testing loss varies from 0.0147 to 0.6586. In Fig. 3e, the training loss converges to the value of  $2.7095e-05$  and the testing loss varies from 0.0092 to 1.2760 throughout the whole simulation process.

Fig. 4 shows the training history for the global model. Here, the global model is created with the aggregation of local models and the evaluation results are conducted based on the training data of all clients. After some communication rounds, as the global model has gained reasonable predictive capacity, its training loss becomes stable and satisfied with the value of  $1.4985e-04$  at  $30^{th}$  communication round. During the testing process, as the global model is aggregated by all the local models, it has already gained the predictive capacity based on various local datasets and the testing loss is stable and satisfied with the value of  $4.1726e-03$  at the  $30^{th}$  communication round. According to Fig. 3 and Fig. 4, there is a proof that our proposed federated learning scheme outperforms the centralized approaches in testing with different distributions of clinical data.

The sample prediction results evaluated on the testing dataset are shown in Fig. 5, 6, 7, and 8, where each prediction represents tuberculosis, diabetes, bronchial asthma, and pneumonia, respectively. As a result of utilizing the attention mechanism in sequence-to-sequence network architecture, the

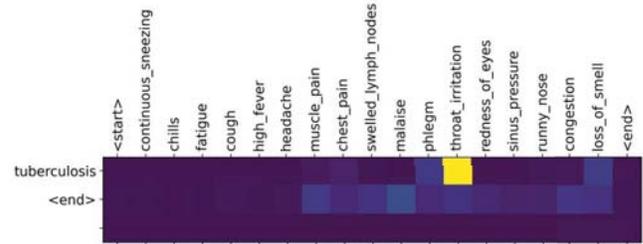


Fig. 5: Prediction of tuberculosis.

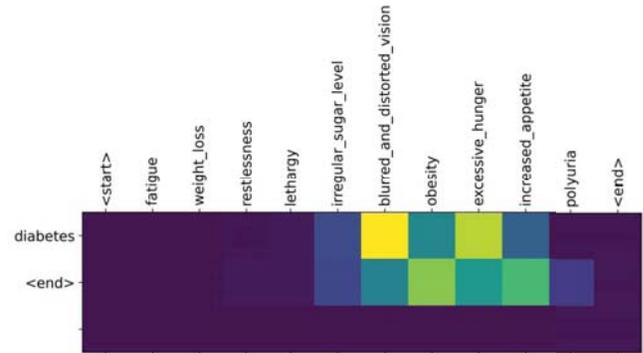


Fig. 6: Prediction of diabetes.

prediction results of diseases are shown with attention weights, where the most relevant symptom is represented by the yellow color.

## V. CONCLUSIONS

In this paper, we proposed a deep learning-based personalized clinical decision support system trained and managed under a federated learning paradigm to assist healthcare professionals in medical diagnosing. We integrated the aid of edge AI to train the local models distributively on the personal

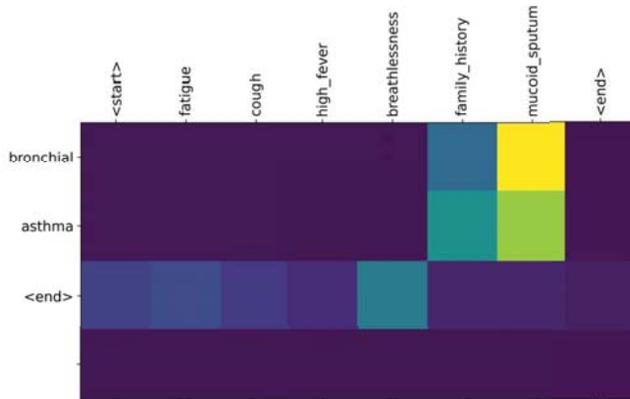


Fig. 7: Prediction of bronchial asthma.

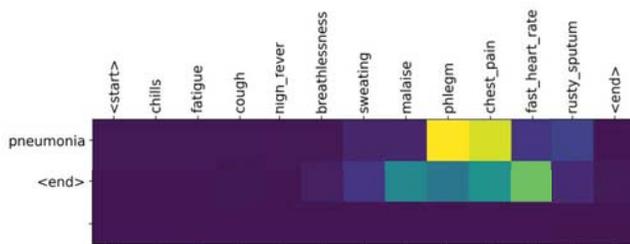


Fig. 8: Prediction of pneumonia.

devices of healthcare professionals for reducing the data communication cost and avoiding the transmission of vulnerable data. We exploited the federated learning framework to train our personalized models collaboratively by learning a shared global model at the central server while keeping all the privacy of clinical data on each local device. As a result, our proposed clinical decision support system outperforms the benchmark centralized scheme by enabling large-scale clinical data mining while preserving the sensitive information of patients and medical organizations. Moreover, we adopted the attention mechanism in designing our medical diagnosis models as sequence-to-sequence networks. Thus, our system possesses the evolvable characteristics that can easily adapt new input symptoms and expand the prediction for new diseases. For a more detailed scenario as optimizing the communication costs of the federated model to account for the noisiness of various biomedical data is considered as our future work.

## REFERENCES

- [1] A. Zanella, N. Bui, A. Castellani, L. Vangelista, and M. Zorzi, "Internet of things for smart cities," *IEEE Internet of Things Journal*, vol. 1, no. 1, pp. 22–32, feb 2014.
- [2] C. M. Thwal, K. Thar, and C. S. Hong, "Edge ai based waste management system for smart city," *Journal of the Korean Information Science Society*, pp. 403–405, Jun. 2019. [Online]. Available: <http://www.dbpia.co.kr/journal/articleDetail?nodeId=NODE08763198>
- [3] C. M. Thwal and C. S. Hong, "A uav-assisted intelligent delivery system for smart city," *Journal of the Korean Information Science Society*, pp. 314–316, Dec. 2019. [Online]. Available: <http://www.dbpia.co.kr/journal/articleDetail?nodeId=NODE09301572>

- [4] Y. L. Tun and C. S. Hong, "An edge-based vehicle surveillance system for enforcing vehicle restriction policies," *Journal of the Korean Information Science Society*, pp. 341–343, Dec. 2019. [Online]. Available: <http://www.dbpia.co.kr/journal/articleDetail?nodeId=NODE09301581>
- [5] S. P. Mohanty, U. Choppali, and E. Kougiianos, "Everything you wanted to know about smart cities: The internet of things is the backbone," *IEEE Consumer Electronics Magazine*, vol. 5, no. 3, pp. 60–70, jul 2016.
- [6] "Disease outbreak news (dons)," Sep 2020. [Online]. Available: <https://www.who.int/csr/don/en/>
- [7] "Preventing epidemics and pandemics." [Online]. Available: <https://www.who.int/activities/preventing-epidemics-and-pandemics>
- [8] J. Liu, X. Kong, F. Xia, X. Bai, L. Wang, Q. Qing, and I. Lee, "Artificial intelligence in the 21st century," *IEEE Access*, vol. 6, pp. 34 403–34 421, 2018.
- [9] V. Carchiolo, A. Longheu, G. Reitano, and L. Zagarella, "Medical prescription classification: a NLP-based approach," in *Proceedings of the 2019 Federated Conference on Computer Science and Information Systems*. IEEE, sep 2019.
- [10] S. Dash, S. K. Shakyawar, M. Sharma, and S. Kaushik, "Big data in healthcare: management, analysis and future prospects," *Journal of Big Data*, vol. 6, no. 1, jun 2019.
- [11] J. Xu, B. S. Glicksberg, C. Su, P. Walker, J. Bian, and F. Wang, "Federated learning for healthcare informatics," 2019.
- [12] Y.-L. Lee, P.-K. Tsung, and M. Wu, "Technology trend of edge AI," in *2018 International Symposium on VLSI Design, Automation and Test (VLSI-DAT)*. IEEE, apr 2018.
- [13] T. Li, A. K. Sahu, A. Talwalkar, and V. Smith, "Federated learning: Challenges, methods, and future directions," *IEEE Signal Processing Magazine*, vol. 37, no. 3, pp. 50–60, may 2020.
- [14] T. S. Brisimi, R. Chen, T. Mela, A. Olshesky, I. C. Paschalidis, and W. Shi, "Federated learning of predictive models from federated electronic health records," *International Journal of Medical Informatics*, vol. 112, pp. 59–67, apr 2018.
- [15] O. Choudhury, A. Gkoulalas-Divanis, T. Salonidis, I. Sylla, Y. Park, G. Hsu, and A. Das, "Differential privacy-enabled federated learning for sensitive health data," 2019.
- [16] T. Luong, H. Pham, and C. D. Manning, "Effective approaches to attention-based neural machine translation," in *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, 2015.
- [17] K. Thar, K. T. Kim, Y. L. Tun, C. M. Thwal, and C. S. Hong, "Evolvable symptom-disease investigator for smart healthcare decision support system," *Journal of the Korean Information Science Society*, pp. 578–580, Jul. 2020. [Online]. Available: <http://www.dbpia.co.kr/journal/articleDetail?nodeId=NODE09874510>
- [18] D. Bahdanau, K. Cho, and Y. Bengio, "Neural machine translation by jointly learning to align and translate," 2014.
- [19] M. Xie, G. Long, T. Shen, T. Zhou, X. Wang, and J. Jiang, "Multi-center federated learning," 2020.
- [20] Aniruddha-Tapas, "Aniruddha-tapas/predicting-diseases-from-symptoms," Apr 2017. [Online]. Available: <https://github.com/Aniruddha-Tapas/Predicting-Diseases-From-Symptoms/tree/master/Manual-Data>
- [21] X. Wang, A. Chused, N. Elhadad, C. Friedman, and M. Markatou, "Automated knowledge acquisition from clinical narrative reports," *Journal of the American Medical Informatics Association*, pp. 783–787, Nov. 2008.
- [22] C. Friedman, L. Shagina, Y. Lussier, and G. Hripcsak, "Automated encoding of clinical documents based on natural language processing," *Journal of the American Medical Informatics Association*, vol. 11, no. 5, pp. 392–402, sep 2004.
- [23] D. Bahdanau, J. Chorowski, D. Serdyuk, P. Brakel, and Y. Bengio, "End-to-end attention-based large vocabulary speech recognition," *2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, Mar 2016. [Online]. Available: <http://dx.doi.org/10.1109/ICASSP.2016.7472618>
- [24] M. Abadi, A. Agarwal, P. Barham, E. Brevdo, Z. Chen, C. Citro, G. S. Corrado, A. Davis, J. Dean, M. Devin, S. Ghemawat, I. Goodfellow, A. Harp, G. Irving, M. Isard, Y. Jia, R. Jozefowicz, L. Kaiser, M. Kudlur, J. Levenberg, D. Mane, R. Monga, S. Moore, D. Murray, C. Olah, M. Schuster, J. Shlens, B. Steiner, I. Sutskever, K. Talwar, P. Tucker, V. Vanhoucke, V. Vasudevan, F. Viegas, O. Vinyals, P. Warden, M. Watteberg, M. Wicke, Y. Yu, and X. Zheng, "Tensorflow: Large-scale machine learning on heterogeneous distributed systems," 2016.