






# Federated Health Recommender System

Sarah Pinon<sup>1</sup><sup>a</sup>, Simon Jacquet<sup>1</sup>, Colin Vanden Bulcke<sup>2</sup><sup>b</sup>, Edouard Chatzopoulos<sup>3</sup><sup>c</sup>,  
Xavier Lessage<sup>4</sup><sup>d</sup> and Raphaël Michel<sup>4</sup><sup>e</sup>

<sup>1</sup>*NaDI, Namur Digital Institute, University of Namur, Belgium*

<sup>2</sup>*IONS, Institute of NeuroScience, UCLouvain, Belgium*

<sup>3</sup>*ICTEAM, Institute of Information and Communication Technologies, Electronics and Applied Mathematics, UCLouvain, Belgium*

<sup>4</sup>*CETIC, DSIDE, Charleroi, Belgium*

**Keywords:** Health, Precision Medicine, Patient-Oriented Decision Support System, Recommender System, Federated Learning.

**Abstract:** Precision Medicine is a new and growing approach to health care. This initiative includes different patient-oriented Decision Support Systems (DSS), such as Health Recommender Systems (HRS). These patient-oriented DSS aim to increase the accuracy and personalization of health care. However, the development of these systems faces a major obstacle related to the confidential and private nature of medical data. These systems require, indeed, a large volume of data to run effectively. But medical data are dispersed among several institutions and cannot be centralized for strict confidentiality reasons. To address this issue, this position paper proposes a system's architecture in which Federated Learning is exploited to build a HRS. Federated Learning allows exploiting the data maintained by different institutions to build the system without requiring their sharing. To demonstrate the feasibility of our proposition, we build a Federated Drug Recommender System. The goal of the system is to assist doctors in their administration of drugs by using historical disease-drug interactions and drug data. As a position paper, the objective of this use case is limited to a proof of concept realized on non-sensitive open-source data. Our ambition is then to use the architecture proposed in this paper to develop a Federated HRS on real medical data.

## 1 INTRODUCTION


Precision Medicine is currently increasing in popularity with its ambitious promises of achieving an incomparable quality of healthcare (Koenig et al., 2017; Kosorok and Laber, 2019). To achieve its promises, this healthcare approach aims to offer personalized treatment strategies to patients by considering a large amount of data. This data includes, for example, the patient's clinical condition, lifestyle, genetics, biomarker information. By including all this information, this medicine is able to offer more precise therapeutic strategies (Koenig et al., 2017).


The aspiration of Precision Medicine is perfectly


aligned with the mission of Health Recommendation Systems (HRS). HRS are, indeed, intended to help healthcare professionals to deal with medical information overload (e.g. medical results, drug information, pathology information) in their patient decision-making (Tran et al., 2021). These patient-oriented decision support systems consider these large volumes of data to recommend, for example, the most appropriate medication, treatment or lifestyle (De Croon et al., 2021).


The performance of HRS, like any RS, depends crucially and positively on the volume of data exploited (Tan et al., 2020). In the era of Big Data, the volume of medical data available is increasingly important, coming from clinical institutions, individual patients, pharmaceutical industries,...(Xu et al., 2021). This emergence represents a great opportunity for HRS.


However, the accessibility of this data is challeng-

<sup>a</sup> <https://orcid.org/0000-0001-5392-1020>

<sup>b</sup> <https://orcid.org/0000-0002-2210-1625>

<sup>c</sup> <https://orcid.org/0000-0001-5848-7999>

<sup>d</sup> <https://orcid.org/0000-0003-0861-0068>

<sup>e</sup> <https://orcid.org/0000-0001-5505-9171>

ing. Indeed, medical data is, in the hospital context for example, fragmented across several hospitals. For confidentiality and privacy reasons, these data cannot be easily accessed and centralized on a server outside the institutions to build an effective RS. This limits the creation of a robust, generalizable and high-quality system. Conversely, by developing small, local systems within each hospital, several biases will impact the quality of the results, such as considering only a subset of conditions or treatments (Xu et al., 2021; Rieke et al., 2020).

To exploit medical data from several hospitals to build a generalizable model, we propose to use Federated learning (FL). The privacy and security concerns very important in the medical sector are, by this way, addressed. Indeed, the FL allows institutions to train the RS on their own infrastructure without exchanging raw (e.i. sensitive) data. The only information shared between the different parties for the construction of the global model corresponds to the model parameters (Kaissis et al., 2020; Li et al., 2020; Rieke et al., 2020; Xu et al., 2021).

In this paper, we propose the architecture of a HRS exploiting the FL, named Federated Health Recommender System (F-HRS). The particularity of the proposed architecture lies in the integration of FL to build the recommendation model. In this way, our system overcomes the medical data confidentiality issues stated above. This position paper demonstrates then the feasibility of our proposition by using non-sensitive and open-source data to develop a Federated Drug Recommender System. In future work, we will exploit real data from hospitals to demonstrate the usefulness and effectiveness of our system in real situations.

This paper is organized as follows: Section 2 reviews the state of the literature about HRS and FRS. Section 3 develops the architecture of the proposed system, the selected model used in our system and the federated learning method exploited. Section 4 presents our proof of feasibility. Section 4 discusses our future works. Section 5 concludes the paper.

## 2 LITERATURE REVIEW

Recommendation systems in the health field are intended for two types of users: (1) patients or healthy users and (2) health professionals. In the case of patients, these RS allow to involve patients in the co-creation of their own health. HRS for healthcare professionals aim at assisting these professionals to provide precise patient-oriented decision (De Croon et al., 2021; Tran et al., 2021). In this project, we

focus on HRS for healthcare professionals. The majority of these HRS concern, in the literature, the recommendation of drugs for different pathologies (e.g. diabetes, migraine, infectious diseases).

A major challenge faced by existing HRS is the accessibility of medical data that is sensitive and confidential (Tran et al., 2021). To deal with this problem, the commonly used approach is data encryption. This method allows to keep the confidentiality of the data while exploiting them in the system of recommendations (Tran et al., 2021). Indeed, data encryption, such as homomorphic encryption, involves encryption of the data by the users before sending it to another party. It is then the encrypted data, and not the raw and complete data, which is used to compute the interest function (Hoens et al., 2013). However, this approach of data encryption raises problems due to its high computational and communication costs, which considerably reduces the performance of the recommendation system (Tran et al., 2021).

A technique increasingly used by RS to exploit larger and more widely distributed data while preserving its confidentiality and security is Federated Learning (Alamgir et al., 2022). This type of learning has the particularity, as opposed to data encryption, to train the algorithm where the data are located (e.g. in each hospital). After this local training, the model parameters are, after an encryption, returned to the central model. By this way, the data confidentiality and privacy are protected which is especially important in the medical field (Kaissis et al., 2020).

To the best of our knowledge, despite the interest of Federated Learning for Recommendation Systems and medical data accessibility issues, no HRS has yet integrated the Federated Approach to exploit all medical data while preserving their privacy.

## 3 F-HRS ARCHITECTURE

F-HRS architecture proposed has three main components, namely: the Input Data, the Recommender Algorithm and the Federated Method. This section develops each of these components. Figure 1 illustrates the combination of these components in an architectural way. This illustration is further explained in the following subsections.

### 3.1 Input Data

F-HRS exploits different types of databases, as shown in Figure 1. On the one hand, our system uses hospital data related to patients such as their medical history, their profile, the results of their medical examinations.

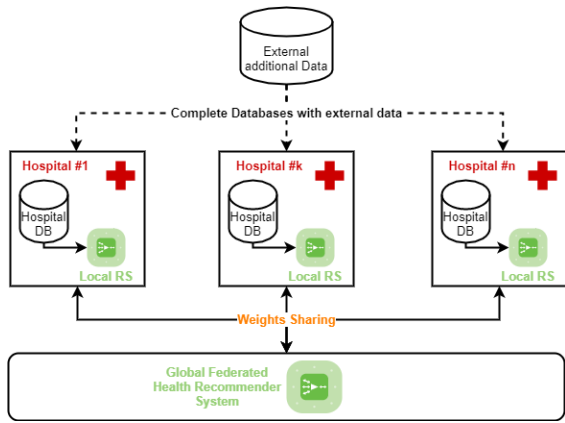


Figure 1: F-HRS Architecture.

These data are in the database of the individual hospitals and cannot be shared due to their private status. On the other hand, our system can be fed by an external database containing information on the recommendation items. This database, unlike the others, is open-source and can therefore be used by the different hospitals.

### 3.2 Recommender Algorithm

The Recommender Algorithm used by our F-HRS is the Neural Collaborative Filtering (NCF) Algorithm (He et al., 2017a). This algorithm works on the principle of collaborative filtering. This filtering is famous in RS and consists in finding users similar to the user of interest and recommending to this user items liked or used by these similar users (Schafer et al., 2007). We have chosen this algorithm because the integration of a neural network within the Collaborative Filtering Algorithm allows to learn the interaction function between users and items automatically from the data rather than manually as is the case in the basic algorithm. Moreover, the NCF Algorithm allows to easily integrate content features representing users and/or items and therefore offer more accurate and relevant recommendations (He et al., 2017b).

Concretely, the NCF algorithm used is illustrated in Figure 2. To process, the algorithm requires, as inputs, positive and negative examples of interactions between users and items. In our case, we applied the general approach which consists of randomly generate four negative examples for each user or patient by using items for which the user has not yet interacted. Positive and negative interactions between users and items are stored into binary matrix. The users or items matrix can, eventually, be enriched with additional data, by concatenating the binary vectors, as shown in Figure 2 for the items case. Each complete user's vector and item's vector from the matrix are then em-

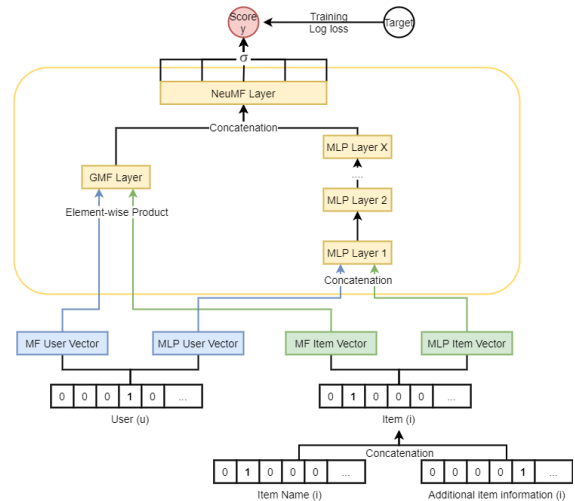


Figure 2: Neural Collaborative Filtering Architecture.

bedded in  $d$  dimensions, eight by default. These vectors are concatenated and introduced in the neural network. Inputs are passed to the hidden layers of the multilayer neural network to learn these non-linear relationships between users and items. Most NCF implementations use 4-5 dense layers with less than 100 neurons per layer. The last layer of the network is a sigmoid function which maps its inputs to a score between zero and one to represent the probability that the user interacts with the item. The algorithm is trained on five epochs, which is a standard. For each epoch, a new set of negative examples is composed for each user.

### 3.3 Federated Learning

To integrate the Federated approach to the HRS, we follow the five general steps, namely (Nishio and Yonetani, 2019):

1. **Initialization:** A global model is initialized randomly or by pre-training with public data. For our experiments, we randomly initialized the weights, which is a common practice in federated learning.
2. **Client Selection:** The different parties (e.i. hospitals) which will contribute to the federated training are selected. By default, our framework selects clients (e.i. hospital servers) in their start order without taking into account a possible client to be discarded. If the need arises, it is quite possible to add additional functionality to this framework to select clients based on various criteria (e. i. such as performance or dataset quality).
3. **Distribution:** The weights from the global model are shared with the parties. This step is critical and must incorporate the best approach to se-

cure weight distribution (P. Treleaven and Pithadia, 2022) based on, for example, homomorphic encryption (HE) classical encryption (TLS), differential privacy (DP) or Secure Multi-party Computation (SMPC).

4. **Update and Upload:** Each party (e.i. hospital) trains the model (e.i. Collaborative Filtering Algorithm) on their local data, and subsequently uploads the updates model weights to the server.
5. **Aggregation:** The weights collected are aggregated and the global model is replaced with the model resulting from the aggregation. The aggregation can be carried out by means of a weighted sum for example, but there are many methods depending on the type of data to be trained. In our system, we use the Federated Averaging method. This aggregation method is the most prevalent one, because of both its simplicity and its performance. Federated averaging consists in computing the weight of the global model as the average of the weights of the clients models.

All steps except the **Initialization** are iterated until the global model achieves a desired level of performance.

Various frameworks are available for the deployment of this Federated approach. In the context of this project, we opted for Flower (Beutel et al., 2021). Flower is an open source framework dedicated to facilitate the deployment of machine learning algorithms in federated environments. The following figure 3 shows its architecture based on an aggregation server and several clients running with different environments. The main particularity of Flower compared to other Federated Learning frameworks is its support for a wide variety of machine learning back-end libraries such as Pytorch, Tensorflow or Jax thanks to its ML framework-agnostic implementation. Flower is also flexible, as it allows both single host simulation and multi-host deployment of the federated procedure (Liu et al., 2022).

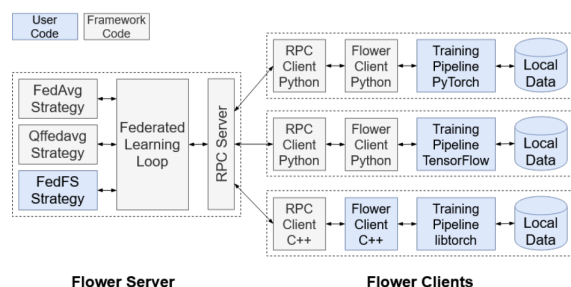


Figure 3: Flower core framework architecture.

## 4 PROOF OF FEASIBILITY

To demonstrate the feasibility of the proposed architecture, we develop a simple Federated Drug Recommender System based on non-sensitive open-source data. To goal of this Recommender System is to assist the doctors in the drugs administration.

To achieve this, our system exploits two types of databases. The first concerns interactions between drugs and diseases (Dat, 2018). The second is related to drugs' features (e.g. pharmaceutical class, contraindications) (Dru, 2022).

To build our system, we divide the data set with the drugs pathology's interactions in a training and test sets. The test set contains for each pathology the last drug related in the database. First, we used the training set to build our system following the architecture developed in the previous sections. In a second step, we evaluated the performance of our system by checking if the last drug associated to the pathology for which we search for a drug is in the ten drugs recommendations made by our system.

The performance of our Federated Drug Recommender System is not presented in this paper. These figures are, indeed, at the moment irrelevant since our system is still being improved and requires better and real data to illustrate the relevance of the approach. Nevertheless, thanks to our use case, the technical feasibility of the architecture could be validated.

## 5 POTENTIAL ISSUES AND CONCERNS

During the development of our system architecture, we considered potential problems related to our application domain (i.e. the medical sector) and to the applied approach.

First, the medical field is facing a lack of standardization of data formats and protocols which can negatively affect the integration of data from multiple sources. To face this problem, we have two solutions. In a first step, it is possible to focus on data collected by widely used software and therefore store it according to a certain standardized structure. Secondly, this problem of standardization is known within the literature and the medical section. It generates, indeed, the development of standardized data formats such as OMOP (Klann et al., 2019; Makadia and Ryan, 2014). These new data formats would solve this heterogeneity problem.

Second, implementing the federated approach is more resource intensive than a traditional centralized system. Nevertheless, the potential benefits of a fed-

erated approach in the medical field can far outweigh the challenges of implementing it.

Third, the risk of data bias and incorrect recommendations to the incomplete data in any machine learning system is a reality. With the federated learning, we try to reduce that risk by training a module with more data coming from multiple sources. By this way, more comprehensive data can be available but also more diverse datasets. Non-IID data across data providers can, indeed, be a major issue in FL. However, a lot of research (Li et al., 2019; Zhao et al., 2018) is carried out and techniques are being developed to improve the quality of federated systems in such settings.

## 6 FUTURE WORK

Having developed the architecture of our F-HRS and demonstrated its technical feasibility, our next goal is to operationalize this system on real data. To achieve this, we are in discussion with different hospitals to get our local models developed on their real database. This step will obviously require some parametric adaptations of the system.

Moreover, after several discussions with health care professionals, it is clear that hospital data is rarely in a structured format. Generally, medical reports are available as PDF documents written in natural language. To be able to integrate this important information into our recommender system, an important step related to the data preparation must be considered. To do so, different techniques of Natural Language Processing will have to be mobilized in order to transform these unstructured data into structured data.

## 7 CONCLUSIONS

In this position paper, we propose to integrate the Federated Learning to build a health recommender model. By this way, we want to overcome the lack of data that medical institutions face in developing patient-oriented decision support systems. Indeed, federated approach allows to train a general and robust recommender model on data from several institutions without the need to share the raw data. The technical feasibility of our solution has been demonstrated via open-source data in the context of a drug-recommender system.

## ACKNOWLEDGEMENTS

We would like to thank the Walloon region for the funding of the ARIAC project, of which our F-HRS project was born and is part. Thanks also to the TRAIL organization which organized the TRAIL Summer Workshop in September 2022 during which our F-HRS project was developed.

## REFERENCES

- (2018). Recommendation medicines by using a review. <https://www.kaggle.com/code/choccozz/recommendation-medicines-by-using-a-review/data>. Accessed on 2022-09-08.
- (2022). Drugcentral. <https://drugcentral.org/>. Accessed on 2022-09-08.
- Alamgir, Z., Khan, F. K., and Karim, S. (2022). Federated recommenders: methods, challenges and future. *Cluster Computing*, pages 1–22.
- Beutel, Qiu, T. M., Parcollet, de Gusmao, and Lane (2021). Flower: A friendly federated learning framework. *On-device Intelligence Workshop at the Fourth Conference on Machine Learning and Systems (MLSys)*.
- De Croon, R., Van Houdt, L., Htun, N. N., Štiglic, G., Abeele, V. V., Verbert, K., et al. (2021). Health recommender systems: systematic review. *Journal of Medical Internet Research*, 23(6):e18035.
- He, X., Liao, L., Zhang, H., Nie, L., Hu, X., and Chua, T.-S. (2017a). Neural collaborative filtering. <https://arxiv.org/>.
- He, X., Liao, L., Zhang, H., Nie, L., Hu, X., and Chua, T.-S. (2017b). Neural collaborative filtering. In *Proceedings of the 26th international conference on world wide web*, pages 173–182.
- Hoens, T. R., Blanton, M., Steele, A., and Chawla, N. V. (2013). Reliable medical recommendation systems with patient privacy. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 4(4):1–31.
- Kaissis, G. A., Makowski, M. R., Rückert, D., and Braren, R. F. (2020). Secure, privacy-preserving and federated machine learning in medical imaging. *Nature Machine Intelligence*, 2(6):305–311.
- Klann, J. G., Joss, M. A., Embree, K., and Murphy, S. N. (2019). Data model harmonization for the all of us research program: Transforming i2b2 data into the omop common data model. *PLoS one*, 14(2):e0212463.
- Koenig, I. R., Fuchs, O., Hansen, G., von Mutius, E., and Kopp, M. V. (2017). What is precision medicine? *European respiratory journal*, 50(4).
- Kosorok, M. R. and Laber, E. B. (2019). Precision medicine. *Annual review of statistics and its application*, 6:263.
- Li, T., Sahu, A. K., Talwalkar, A., and Smith, V. (2020). Federated learning: Challenges, methods, and future directions. *IEEE Signal Processing Magazine*, 37(3):50–60.

- Li, X., Huang, K., Yang, W., Wang, S., and Zhang, Z. (2019). On the convergence of fedavg on non-iid data. *arXiv preprint arXiv:1907.02189*.
- Liu, X., Shi, T., Xie, C., Li, Q., Hu, K., Kim, H., Xu, X., Li, B., and Song, D. X. (2022). Unifed: A benchmark for federated learning frameworks. *ArXiv*, abs/2207.10308.
- Makadia, R. and Ryan, P. B. (2014). Transforming the premier perspective® hospital database into the observational medical outcomes partnership (omop) common data model. *Egems*, 2(1).
- Nishio, T. and Yonetani, R. (2019). Client selection for federated learning with heterogeneous resources in mobile edge. *IEEE International Conference on Communications*, 2019-May.
- P. Treleaven, M. S. and Pithadia, H. (2022). Unifed: Federated learning: The pioneering distributed machine learning and privacy-preserving data technology. *iee*.
- Rieke, N., Hancox, J., Li, W., Milletari, F., Roth, H. R., Albarqouni, S., Bakas, S., Galtier, M. N., Landman, B. A., Maier-Hein, K., et al. (2020). The future of digital health with federated learning. *NPJ digital medicine*, 3(1):1–7.
- Schafer, J. B., Frankowski, D., Herlocker, J., and Sen, S. (2007). Collaborative filtering recommender systems. In *The adaptive web*, pages 291–324. Springer.
- Tan, B., Liu, B., Zheng, V., and Yang, Q. (2020). A federated recommender system for online services. In *Fourteenth ACM Conference on Recommender Systems*, pages 579–581.
- Tran, T. N. T., Felfernig, A., Trattner, C., and Holzinger, A. (2021). Recommender systems in the healthcare domain: state-of-the-art and research issues. *Journal of Intelligent Information Systems*, 57(1):171–201.
- Xu, J., Glicksberg, B. S., Su, C., Walker, P., Bian, J., and Wang, F. (2021). Federated learning for healthcare informatics. *Journal of Healthcare Informatics Research*, 5(1):1–19.
- Zhao, Y., Li, M., Lai, L., Suda, N., Civin, D., and Chandra, V. (2018). Federated learning with non-iid data. *arXiv preprint arXiv:1806.00582*.