THESIS project

Federated machine learning for healthcare applications based on medical imaging

Summary

Application of machine learning (ML) to healthcare is among the most challenging ones with the potential to exploit information provided by an exponentially growing mass of heterogeneous data (images, semantic information, biological parameters,..). Those models require a large amount of data to perform well, particularly in the era of large-scale deep neural networks. One option to increase the training population is to promote multi-centre clinical studies, which opens many privacy-related problems since data producers lose control over their data as well as huge data traffic. Federated learning (FL) is a new ML approach that was recently introduced to counterbalance the need to access large databases by the responsibility to maintain the privacy of individual participants. In this context, FL appears as a very promising technique, first to account for patient privacy thus complying with the increasingly stringent general data protection regulations (GDPR) and then to limit the huge amount of data traffic required when gathering medical data to a centralized server. This research field is in its early premise and needs to address key challenges related to the specificity of medical data. The aim of this PhD is to investigate methodological research in this domain with application to the design of diagnosis and prognosis models of brain pathology based on multimodality imaging.

Keywords : Federated Machine Learning – Medical Imaging – Diagnosis and Prognosis Models – Privacy and Security

Thesis supervision and collaboration

The PhD candidate will be co-supervised by Stefan Duffner (MCF INSA HDR – LIRIS) and Carole Lartizien (DR CNRS – CREATIS). **Carole Lartizien**, has expertise in machine learning for medical image analysis. She will contribute to provide her expertise on AI based diagnosis models of neuroimaging data and access to the imaging datasets. **Stefan Duffner** has strong expertise in machine learning for computer vision. He will contribute to this collaboration with his extensive knowledge and competence in neural network models and learning algorithms under challenging conditions (e.g. few, non i.i.d. data; uncertainty/noise; heterogeneity) applied to this new medical image analysis context.

Working environment and salary

The doctoral student will share his/her time between the two laboratories according to the needs and the progress of the project. He/she will participate in team meetings in both laboratories to benefit as much as possible from the two scientific stimulating environment of the MYRIAD and IMAGINE teams. The PhD candidate will benefit from ongoing collaborations with external experts in the machine learning and neuroimaging domains. He/she will have access to multimodality imaging databases that have been collected for the two use case applications considered in the PhD project.

Employment would ideally start in Fall 2021 and is funded half by INSA Lyon and the French National Research Agency (ANR) under the IADoc Research program. Salary is around 1700 euros net per month (+ teaching)

Profil of the Applicant

The candidate is expected to have strong knowledge in machine learning and some experiment in image processing. Some prior experience with medical image processing would be appreciated but is not required. Good programming skills (python..) are also required. We are looking for an enthusiastic and autonomous student with strong motivation and interest in multidisciplinary research (image processing and machine learning in a medical context).

Contact

For more details on the position, please contact carole.lartizien@creatis.insa-lyon.fr and Stefan.duffner@insa-lyon.fr with your CV.

The formal application procedure will be detailed upon request. Deadline for application is **23rd of April 21**.

Detailed Thesis project

Federated machine learning for healthcare applications based on medical imaging

Context

Traditional machine learning (ML) models are trained on a central server, which involves centralizing the data created on edge devices. This opens many privacy-related problems since users lose control over their data. The data flow generated can also be a limiting factor. In order to address these issues, a new ML approach was recently introduced called federated learning (FL) where copies of the current model are locally trained and updated on each device using local data, then sent back to the central server where it is aggregated into a single improved global model which is sent back to the devices.

Application of ML to healthcare is among the most challenging ones with the potential to exploit information provided by exponentially growing mass of heterogeneous data (image, semantic information, biological parameters etc..). In this context, FL appears as a very promising technique, first to account for patient privacy thus complying with the increasingly stringent general data protection regulations (GDPR) and then to limit the huge amount of data traffic required when gathering medical data to a centralized server. This research field is in its early premise and needs to address key challenges related to the specificity of medical data.

Research program

The aim of this PhD is to investigate methodological research in the field of federated learning for medical image analysis, and more specifically for the design of diagnosis and prognosis models of brain pathology based on multimodality imaging.

Medical diagnosis or prognosis models are designed to assist clinicians either by highlighting abnormal regions in an image, predicting a diagnosis or patient outcome. Those models require a large amount of data to perform well, particularly in the era of large-scale deep neural networks. One option to increase the training population is to promote multi-center clinical studies, which allow gathering small to medium size heterogeneous datasets located in different clinical centers. In this context, federated learning is extremely appealing to counterbalance the need to access large patient cohorts by the responsibility to maintain the privacy of individual participants.

Current FL approaches generally distribute copies of a machine learning algorithm to the sites or devices where the data is kept (nodes), performing training iterations locally, and returning the results of the computation (for example, updated neural network weights) to a central repository to update the main algorithm. However, simple distributed training does not offer provable privacy guarantees to satisfy technical safe standards and may reveal information about the underlying patients. Moreover, training more complex models, such as deep neural networks, distributed over many sites may considerably increase the number and volume of exchanged messages and traffic and entail scalability and security issues. Therefore, more effective algorithms of distributed training of these models need to be developed, and the structure of these models (e.g. neural network architecture) needs to be adapted to this distributed scenario. Finally, the current trend in machine learning for medical image analysis is to develop non-conventional deep architectures accounting for the specificity of medical data, e.g. learning with weak, partial, uncertain annotations, heterogeneous data (multi-center). The embedding of such complex architectures within the framework of FL is thus another key challenge to address.

This leads to several fundamental scientific challenges and questions in distributed and privacy-preserving learning such as :

- What is the optimal method for updating the central model state (distributed optimization, federated averaging)?
- What is the influence of training different (distributed) models on the convergence and efficiency as well as on test performance?
- Should we consider one monolithical ML model or a more modular architecture or another organisation potentially reducing the communication between servers and increasing security and privacy?
- How to quantify the robustness to heterogeneous, imbalanced data and complex deep architectures? The decentralized nature of the data, as an example, complicates data curation to ascertain the integrity and quality of the results.
- How to combine the distributed training with other methods to guarantee security and privacy (e.g., using trusted execution environments to secure the backend server, strongly encrypted messages)?

CREATIS will provide different use case diagnosis models of brain pathologies based on multi-modality (MRI, PET) neuroimaging data :

- We will start with a deep diagnosis model of Alzheimer patients assessing whether their neuroimaging exams present patterns of normal elderly (> 65 y) control subjects (NC), mild cognitive impairment (MCI) or Alzheimer disease (AD). This model will be trained on datasets retrieved from the publicly available ADNI database (Alzheimer's Disease Neuroimaging Initiative) concatenating more than 1500 multimodality exams from more than 50 north American clinical centers (thus including highly heterogeneous data). This use case is considered as a “simple” classification task that can be performed with light and simple deep convolutional neural networks (CNN) architectures. It will allow setting the grounds of the methodological and technological developments and provide preliminary answers on most of the methodological questions listed above.

    - The second model that we will consider is a prototype deep model that we designed to detect brain anomalies in multimodality imaging (MRI, PET) (see [Alaverdyan et al, 2020] in C. Lartizien’s CV). This model is trained on populations of normal healthy subjects. Collecting such large dataset of data from normal subjects is difficult, our aim is thus to gather these data from multiple clinical centers. One of the current applications of this model is the detection of epileptogenic lesions, in collaboration with clinicians from the Hospices Civils de Lyon (HCL). This difficult diagnostic task requires the development of deep architectures more complex than standard CNNs. This application is therefore a more complex use case than the first one and will allow to evaluate and update the pipeline developed on the first application.

The collected data for the two use cases will be first stored in a private and secure storage warehouse from CREATIS. We will emulate a platform that will mimic a network of distant and local sources with secure access to implement and evaluate the FL solutions that will be designed. This infrastructure will simulate hospital radiological PACS (Picture Archiving and Communication System) on virtual machines, modelling heterogeneous architectures and communication protocols.

Expected results

By the end of this PhD project, we expect to have produced (i) top ranked publications; (ii) original open code validated on the two use case medical applications; (iii) a federated learning infrastructure and test in the lab environment.

The work will aim for fundamental contributions in statistical deep learning as well as medical imaging analysis with high-level publications in both communities (NeurIPS, ICML, AAAI, MICCAI, IPMI…)

This PhD would be the opportunity to start a new collaboration in this emerging field of FL for medical imaging and strengthen the link between the two local laboratories. We expect that this PhD work will enable accelerating research for the design of diagnosis and prognosis tools based on medical image analysis by giving access to massive medical datasets (following the stringent GRPD rules) in a completely secure way (keeping privacy).

Provisional calendar:

The research axes to explore will be defined and prioritized according to prevailing methodological challenges to tackle at the beginning of the PhD project. These challenges will be defined from the review of the state-of-the art bibliography. Hereafter, we provide an attempted draft calendar that will be updated.

Year 1 : Biblio study on - federated, privacy preserving and secure learning and – AI models for medical image analysis. Definition and prioritization of the methodological research axes. Handling of the AD datasets and DL codes for the first medical use case. Implementation and analysis of different architectures of federated learning on the virtual PACS networks. First conference paper writing.

Year 2 : Impact analysis of the different parameters of the pipeline (data heterogeneity of PACS and imaging data..) developed for the first use case. Integration of privacy-preserving elements in the pipeline. Handling of the second-use case datasets and corresponding DL codes. First conference paper writing

Year 3 : Update and performance analysis of the final pipeline in the two use-case applications on the virtual PACS nodes architectures. Performance analysis of the final pipeline on the second use-case application in the lab virtual infrastructure based on distant and heterogeneous emulated hospital PACS. Paper writing. Thesis writing and defence.

Thesis supervision

The PhD candidate will be co-supervised by Stefan Duffner (MCF INSA HDR – LIRIS) and Carole Lartizien (DR CNRS – CREATIS). This thesis will contribute to strengthen the ties between regional academic labs in computer sciences on this challenging topic of artificial intelligence. We will organize regular meetings with research scientists interested in federated learning within the framework of local (Fédération d'Informatique Lyonnaise) as well as national networks (GDR ISIS and MadICs).

Reference list on federated learning (non exhaustive)

H. B. McMahan, E. Moore, D. Ramage, S. Hampson, B. Agüera y Arcas. Communication-Efficient Learning of Deep Networks from Decentralized Data. AISTATS, 2017.

M. Mohri, G. Sivek, and A. T. Suresh. Agnostic federated learning. ICML, 2019.

T. Li, A. K. Sahu, A. Talwalkar and V. Smith. Federated Learning: Challenges, Methods, and Future Directions. in *IEEE Signal Processing Magazine*, vol. 37, no. 3, pp. 50-60, 2020.

A. Reisizadeh, F. Farnia, R. Pedarsani, A. Jadbabaie. Robust Federated Learning: The Case of Affine Distribution Shifts, NeurIPS, 2020.

Kaissis, G.A., Makowski, M.R., Rückert, D. *et al.* Secure, privacy-preserving and federated machine learning in medical imaging. *Nat Mach Intell* **2,** 305–311 (2020). https://doi.org/10.1038/s42256-020-0186-1

Beaulieu-Jones, Brett K. and Yuan, Wi and Finlayson, S. G. and Wu, Z S. Privacy-Preserving Distributed Deep Learning for Clinical Data. Machine Learning for Health (ML4H) Workshop at NeurIPS 2018 arXiv:1811.07216

M. J. Sheller, G. A. Reina, B. Edwards, J. Martin, S. Bakas. Multi-Institutional Deep Learning Modeling Without Sharing Patient Data: A Feasibility Study on Brain Tumor Segmentation, MICCAI Brain Lesion (BrainLes) workshop, 2018