



Federated Learning of Biomedical Data in Multicentric Imaging Studies

Santiago Smith Silva Rincon

► To cite this version:

Santiago Smith Silva Rincon. Federated Learning of Biomedical Data in Multicentric Imaging Studies. Mathematics [math]. UCA, Inria, 2023. English. NNT : . tel-04417044v2

HAL Id: tel-04417044

<https://inria.hal.science/tel-04417044v2>

Submitted on 25 Jan 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution - ShareAlike 4.0 International License



PHD THESIS

Federated Learning of Biomedical Data in Multicentric Imaging Studies

Santiago Smith SILVA RINCON

INRIA, EPIONE Team

Supervised by Marco LORENZI and co-supervised by Barbara BARDONI

Defended on July 18, 2023

Presented to obtain the title of DOCTEUR EN AUTOMATIQUE, TRAITEMENT DU SIGNAL ET DES IMAGES of the UNIVERSITÉ CÔTE D'AZUR.

Jury:

James COLE	Dementia Research Centre, Institute of Neurology, UCL	Reviewer
Jean-François	Univ. Paris-Saclay, CEA, CNRS, NeuroSpin, BAOBAB	Reviewer
Marco LORENZI	INRIA Sophia-Antipolis	Supervisor
Barbara BARDONI	INSERM, CNRS, IPMC, Sophia-Antipolis	Co-supervisor
Jonas RICHIARDI	Lausanne University Hospital & Univ. of Lausann	Invited
Andre Altmann	Centre for Medical Image Computing (CMIC), Dept. of Medical Physics and Biomedical Engineering, UCL	Invited



THÈSE DE DOCTORAT

Apprentissage Fédéré de Données Biomédicales dans les Études d'Imagerie Multicentriques

Santiago Smith SILVA RINCON

INRIA, Équipe EPIONE

Thèse dirigée par Marco LORENZI et co-dirigée par Barbara BARDONI

Soutenue le 18 juillet 2023

Présentée en vue de l'obtention du grade de DOCTEUR EN AUTOMATIQUE, TRAITEMENT
DU SIGNAL ET DES IMAGES de l'UNIVERSITÉ CÔTE D'AZUR.

Devant le jury composé de :

James COLE	Dementia Research Centre, Institute of Neurology, UCL	Rapporteur
Jean-François	Univ. Paris-Saclay, CEA, CNRS, NeuroSpin, BAOBAB	Rapporteur
Marco LORENZI	INRIA Sophia-Antipolis	Directeur de thèse
Barbara BARDONI	INSERM, CNRS, IPMC, Sophia-Antipolis	Co-encadrant
Jonas RICHIARDI	Lausanne University Hospital & Univ. of Lausann	Invité
Andre Altmann	Centre for Medical Image Computing (CMIC), Dept. of Medical Physics and Biomedical Engineering, UCL	Invité

Abstract

In order to gather sufficient sample size and representativity of clinical populations, the multi-centric analysis paradigm is often adopted for statistical and machine learning studies of biomedical data, particularly in the field of neuroimaging. Conventional multi-centric analysis paradigms are based on meta-analysis and mega-analysis, often in conjunction with data harmonization, to account for systematic biases and improve the combined analysis of data from multiple sources. However, while meta-analyses are mainly suited for standard statistical testing only, mega-analyses require centralizing the data, which can undermine data privacy and security. Today, data protection regulations such as the General Data Protection Regulation (GDPR) and the California Consumer Privacy Act (CCPA) impose strict governance on sensitive patient information, significantly limiting researchers' access to such data.

Federated learning (FL) is an alternative paradigm to multi-centric studies enabling multiple parties to train a model collaboratively without sharing sensitive data. FL thus addresses data governance challenges while enhancing patients' data privacy. However, to facilitate real-life applications of FL, a series of challenges must be addressed: i) heterogeneity and generalization due to differences in data distributions and discrepancies across different institutions which can result in biased models that do not generalize properly, ii) occasional requirements of considerable amounts of computational resources that hospitals or institutions may not have, limiting its practicality and, ii) a common framework and infrastructure to put in place real-life applications while fulfilling research and governance demands. This thesis aims to contribute to the evolving landscape of neuroimaging research by investigating the potential of FL to transform the way researchers collaborate and analyze data, ultimately paving the way for more efficient and effective advancements in neuroimaging.

We start by addressing the issue of data heterogeneity in federated learning setups, by introducing two methods, namely "Fed-ComBat" and "federated mixed-effect modeling", which aim to perform data harmonization and modeling on heterogeneous data respectively.

Secondly, we introduce a black-box optimization scheme for FL aiming to improve the optimization process in federated setups. This method is based on gradient-free optimization of a global model, through the collaborative iterative refinement of the

cost function associated with the distributed optimization problem across clients. This approach aims to centralize computational costs and mitigate overfitting issues linked to gradient-descent-based approaches while enabling institutions and hospitals with limited computational resources to participate in federated learning setups while achieving accurate and generalizable models.

Finally, to enable and empower real-life federated applications, we introduce Fed-BioMed as an open-source framework for federated learning in healthcare, aiming to fulfill the need for a common collaborative framework that is also compliant with privacy and ethical standards.

Overall, this thesis comprises methodological and technical contributions that tackle the challenges of data heterogeneity, optimization, and infrastructure in federated learning setups for neuroimaging research, with the ultimate goal of facilitating more efficient and effective advancements in healthcare while preserving patient privacy and data governance.

Keywords: federated learning, healthcare, data protection, GDPR, CCPA, medical imaging, data harmonization, meta-analysis, mega-analysis, Fed-BioMed, Bayesian optimization, random effect models, FedComBat.

Résumé

Afin de rassembler une taille d'échantillon suffisante et une représentativité des populations cliniques, le paradigme de l'analyse multi-centrique est souvent adopté pour les études statistiques et d'apprentissage automatique des données biomédicales, en particulier dans le domaine de la neuroimagerie. Les paradigmes d'analyse multi-centrique conventionnels reposent sur la méta-analyse et la méganalyse, souvent conjointement avec l'harmonisation des données, pour tenir compte des biais systématiques et améliorer l'analyse combinée des données provenant de sources multiples. Cependant, alors que les méta-analyses sont principalement adaptées aux tests statistiques standard, les méganalyses nécessitent une centralisation des données, ce qui peut nuire à la confidentialité et à la sécurité des données. Aujourd'hui, les réglementations sur la protection des données telles que le Règlement Général sur la Protection des Données (RGPD) et la California Consumer Privacy Act (CCPA) imposent une gouvernance stricte sur les informations sensibles des patients, limitant considérablement l'accès des chercheurs à ces données.

L'apprentissage fédéré (FL) est un paradigme alternatif aux études multi-centriques permettant à plusieurs parties de former un modèle en collaboration sans partager de données sensibles. Le FL répond ainsi aux défis de la gouvernance des données tout en améliorant la confidentialité des données des patients. Cependant, pour faciliter les applications réelles du FL, une série de défis doit être relevée : i) l'hétérogénéité et la généralisation en raison des différences dans les distributions de données et les écarts entre les différentes institutions qui peuvent entraîner des modèles biaisés qui ne se généralisent pas correctement, ii) les besoins occasionnels en ressources de calcul considérables que les hôpitaux ou les institutions peuvent ne pas avoir, limitant ainsi sa praticité et, ii) un cadre et une infrastructure communs pour mettre en place des applications réelles tout en répondant aux exigences de recherche et de gouvernance. Cette thèse vise à contribuer au paysage évolutif de la recherche en neuroimagerie en étudiant le potentiel du FL pour transformer la manière dont les chercheurs collaborent et analysent les données, ouvrant ainsi la voie à des avancées plus efficaces et efficientes en neuroimagerie.

Nous commençons par aborder la question de l'hétérogénéité des données dans les configurations d'apprentissage fédéré, en introduisant deux méthodes, à savoir "Fed-ComBat" et "modélisation à effets mixtes fédérée", qui visent à réaliser respectivement l'harmonisation des données et la modélisation sur des données hétérogènes.

Deuxièmement, nous introduisons un schéma d'optimisation boîte noire pour le FL visant à améliorer le processus d'optimisation dans les configurations fédérées. Cette méthode est basée sur l'optimisation sans gradient d'un modèle global, grâce à l'affinement itératif collaboratif de la fonction de coût associée au problème d'optimisation distribué entre les clients. Cette approche vise à centraliser les coûts de calcul et à atténuer les problèmes de surajustement liés aux approches basées sur la descente de gradient, tout en permettant aux institutions et aux hôpitaux disposant de ressources informatiques limitées de participer aux configurations d'apprentissage fédéré tout en obtenant des modèles précis et généralisables.

Enfin, pour permettre et renforcer les applications fédérées dans la vie réelle, nous présentons Fed-BioMed, un cadre open source pour l'apprentissage fédéré dans le domaine de la santé, visant à répondre au besoin d'un cadre collaboratif commun qui est également conforme aux normes de confidentialité et d'éthique.

Dans l'ensemble, cette thèse comprend des contributions méthodologiques et techniques qui abordent les défis de l'hétérogénéité des données, de l'optimisation et de l'infrastructure dans les configurations d'apprentissage fédéré pour la recherche en neuroimagerie, avec pour objectif ultime de faciliter des avancées plus efficaces et efficientes dans les soins de santé tout en préservant la confidentialité des patients et la gouvernance des données.

Mots-clés: apprentissage fédéré, santé, protection des données, RGPD, CCPA, imagerie médicale, harmonisation des données, méta-analyse, méganalyse, Fed-BioMed, optimisation bayésienne, modèles à effets aléatoires, FedComBat.

Financial Support

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 847579 (Marie Skłodowska-Curie Actions), and by the INRIA Sophia Antipolis-Méditerranée "NEF" computation cluster.

Data Use Agreements

Data used in the preparation of this article were obtained from the Alzheimer's Disease Neuroimaging Initiative (ADNI) database (<https://adni.loni.usc.edu/>).

Data used in this article were provided by the Australian Imaging Biomarkers and Lifestyle flagship study of ageing (AIBL) funded by the Commonwealth Scientific and Industrial Research Organisation (CSIRO) which was made available at the ADNI database (<https://adni.loni.usc.edu/aibl-australian-imaging-biomarkers-and-lifestyle-study-of-ageing-18-month-data-now-released/>).

Data used in the preparation of this article were obtained from the Autism Brain Imaging Data Exchange (ABIDE) I database.

Data used in this article were obtained from the Autism Brain Imaging Data Exchange (ABIDE) II database (http://fcon_1000.projects.nitrc.org/indi/abide/abide_II.html).

The AIBL researchers contributed data but did not participate in the analysis or writing of this report. AIBL researchers are listed at <https://www.aibl.csiro.au>.

Data used in the preparation of this article were obtained from the MIRIAD database.

Data used in the preparation of this article were obtained from the Open Access Series of Imaging Studies (OASIS) database.

Data used in the preparation of this article were obtained from the Parkinson's Progression Markers Initiative (PPMI) database (<http://www.ppmi-info.org/>).

Data used in the preparation of this article were obtained from the Pre-symptomatic Evaluation of Experimental or Novel Treatments for Alzheimer's Disease (A4) study.

Contents

1	Introduction	7
1.1	Clinical context	7
1.2	Data Protection Regulations: GDPR and CCPA	8
1.3	Multi-centric studies in neuroimaging	9
1.3.1	Meta-analysis and Mega-analysis	9
1.3.2	Federated Learning	13
1.4	Challenges in Federated Learning and impact on Applications to Healthcare	15
1.4.1	Framework and Infrastructure	15
1.4.2	Communication Efficiency and Resource Utilization	15
1.4.3	Heterogeneity	16
1.4.4	Privacy and Security	17
1.4.5	Model Convergence and Stability	17
1.5	Aims and Structure of the Thesis	17
1.6	Publications	18
1.6.1	Collaborations	18
2	FedComBat: Secure Data Harmonization via Federated Learning	21
2.1	Introduction	22
2.2	Methods	25
2.2.1	Generalized ComBat model	25
2.2.2	Federated ComBat: Fed-ComBat	26
2.2.3	Related works	27
2.3	Materials	30
2.3.1	Synthetic data	30
2.3.2	Brain MRI-Data	32
2.4	Results	34
2.4.1	Synthetic data	34
2.4.2	Brain MRI-data	37
2.5	Limitations and Future Directions	37
2.6	Discussion and Conclusion	41
2.7	Funding	41
2.8	Acknowledgments	42
2.9	Comparing centralized and Fed-ComBat's formulation	43
2.10	Harmonization on synthetic data	43

2.11	Harmonization on real data	44
2.11.1	Evidence of non linear effects on brain phenotypes	44
2.11.2	Residual ComBat effects on real data.	44
2.12	Identifiability of ComBat parameters	47
2.13	CRedit author statement	47
3	Federated black-box Bayesian optimization	49
3.1	Introduction	50
3.2	Overview of Black Box Optimization	51
3.2.1	Bayesian optimization	52
3.2.2	Acquisition functions	53
3.3	Methodology	55
3.3.1	Federated Bayesian Optimization	57
3.4	Results	59
3.4.1	Synthetic Data	59
3.4.2	Brain imaging data	61
3.4.3	Classification	62
3.4.4	Variational autoencoder	63
3.5	Conclusions	63
4	Federated Data Harmonization in Biomedical Research using Mixed Effects Models: A Focus on Conditional Variational Autoencoders	65
4.1	Introduction	65
4.2	Fundamentals of Data Harmonization	66
4.2.1	Significance of Data Harmonization in Biomedical Research . . .	67
4.2.2	Challenges and Issues in Data Harmonization	67
4.3	Methods	68
4.3.1	Federated Harmonization with Conditional Variational Autoencoders (CVAEs)	68
4.3.2	Harmonization	69
4.4	Results	70
4.5	Challenges and Future Directions	71
4.6	Conclusions	75
5	Fed-BioMed: A General Open-source Frontend Framework for Federated Learning in Healthcare	77
5.1	Introduction	78
5.2	Goals and Contributions of Fed-BioMed	79
5.3	Related works	79
5.4	Implementation Overview and Architecture of Fed-BioMed	80
5.4.1	Client (Node) Service	80
5.4.2	Central Node (Network) Service	81
5.4.3	Researcher Service	81

5.4.4	Communication scheme	82
5.4.5	Security	83
5.4.6	Traceability	83
5.5	Related work	83
5.5.1	Non-medical federated learning frameworks	83
5.5.2	Medical oriented federated learning frameworks	86
5.6	Usage	88
5.7	Applications and Case Studies	88
5.7.1	MNIST analysis	88
5.7.2	Brain imaging data analysis	90
5.8	Results	91
5.9	Conclusions	92
5.10	Supplementary Material	93
	Conclusion	95
	Bibliography	97

Introduction

1.1	Clinical context	7
1.2	Data Protection Regulations: GDPR and CCPA	8
1.3	Multi-centric studies in neuroimaging	9
1.3.1	Meta-analysis and Mega-analysis	9
1.3.2	Federated Learning	13
1.4	Challenges in Federated Learning and impact on Applications to Health-care	15
1.4.1	Framework and Infrastructure	15
1.4.2	Communication Efficiency and Resource Utilization	15
1.4.3	Heterogeneity	16
1.4.4	Privacy and Security	17
1.4.5	Model Convergence and Stability	17
1.5	Aims and Structure of the Thesis	17
1.6	Publications	18
1.6.1	Collaborations	18

1.1 Clinical context

The rapid advancements in information technology and digitization of healthcare services have led to an unprecedented growth in the volume and variety of data generated by hospitals and healthcare institutions. Electronic health records (EHRs), medical imaging, wearable devices, and genomic data are just a few examples of the diverse sources contributing to this data explosion [Undavia et al., 2020; Miotto et al., 2018]. The increasing adoption of telemedicine, remote monitoring systems, and mobile health applications has further accelerated the production of digital health information [Wosik et al., 2020; Torous et al., 2018]. These massive and complex datasets, often referred to as "big data," hold the potential to revolutionize medical research, improve patient care, and streamline healthcare operations [Murdoch et al., 2013].

The wealth of information contained in these datasets offers numerous opportunities for data-driven research, including the identification of new biomarkers, optimization of treatment strategies, and development of personalized medicine approaches [Schneeweiss,

2014; Jensen et al., 2012]. Additionally, the analysis of this data can provide valuable insights into healthcare service delivery and quality of care [Rumsfeld et al., 2016]. In particular, in the field of neurosciences, heterogeneity plays a pivotal role in medical studies, especially when dealing with neuroimaging data. The understanding of neurodegenerative diseases, such as Alzheimer's disease or Parkinson's, is highly impacted by such heterogeneity [Webb-Vargas et al., 2017]. This implies that larger and more diverse datasets are crucial for uncovering relevant findings and advancing our knowledge of these complex disorders, ultimately contributing to improved healthcare outcomes and patient care.

However, the vast volumes and complexity of neuroimaging data also pose several challenges, including data storage, management, integration, and analysis [Sun et al., 2013]. The sensitive nature of health information necessitates strict adherence to data protection regulations such as the General Data Protection Regulation (GDPR) [Voigt et al., 2017] and the California Consumer Privacy Act (CCPA) [Solove et al., 2020], which can further complicate data sharing and collaboration efforts. The growing need to harness the full potential of healthcare data while overcoming these challenges has spurred significant interest in the development of novel computational methods and infrastructures tailored to the unique requirements of medical research.

1.2 Data Protection Regulations: GDPR and CCPA

As the volume of healthcare data continues to grow, so does the need for robust data protection and privacy regulations. Health information is often sensitive, and its unauthorized disclosure can have serious consequences for individuals and organizations alike. To address these concerns, governments have enacted strict data protection laws, such as the European Union's General Data Protection Regulation (GDPR) [Voigt et al., 2017] and the United States' California Consumer Privacy Act (CCPA) [Solove et al., 2020].

The GDPR, which came into effect in 2018 as a comprehensive data protection regulation applicable to all EU member states. It aims to unify data privacy laws across Europe and empower individuals with greater control over their personal information. The GDPR imposes stringent requirements on organizations that process personal data, including the need to obtain explicit consent from individuals, the right to data portability, and the right to be forgotten [Voigt et al., 2017]. Additionally, the GDPR mandates that organizations implement appropriate technical and organizational measures to ensure data security and privacy, with significant penalties for non-compliance.

The CCPA, which came into effect in 2020, is a similar data protection regulation enacted by the state of California in the United States. It grants Californian residents several

rights concerning their personal information, such as the right to know what personal data is collected, the right to delete personal information held by businesses, and the right to opt-out of the sale of personal data [Solove et al., 2020]. Although the CCPA is not as extensive as the GDPR, it serves as a significant step toward strengthening data protection and privacy in the United States.

Compliance with data protection regulations such as the GDPR and CCPA is a critical consideration for medical researchers working with healthcare data. It can be particularly challenging in the context of multi-institutional collaborations and large-scale data sharing efforts, where data must be anonymized and de-identified to protect patient privacy while preserving its utility for research purposes [El Emam et al., 2014]. Moreover, these regulations may limit data accessibility, potentially hindering the development and deployment of innovative data-driven solutions in healthcare. As a result, there is a growing interest in exploring alternative approaches, such as federated learning, that enable collaborative research and model development without the need to share raw patient data [Brisimi et al., 2018; Higgins et al., 2019].

In the next section, we discuss the most common methods used for multicentric data analysis, and how federated learning may offer a better solution.

1.3 Multi-centric studies in neuroimaging

To address the challenges and maximize the potential of neuroimaging data, researchers often rely on multicentric studies, which involve the collaboration of multiple institutions and research centers. These studies enable the aggregation of larger and more diverse datasets, thereby increasing the statistical power and generalizability of the findings [Webb-Vargas et al., 2017]. Within the context of multicentric studies, conventional approaches are established such as meta-analysis [Glass, 1976] and mega-analysis [Costafreda, 2009] in order to facilitate collaborative research efforts. Meta-analysis and mega-analysis focus on pooling summary statistics or raw data, respectively, from individual studies. However, these approaches may face limitations due to data privacy and governance concerns. As an alternative, federated learning has been proposed as an alternative paradigm that could be extended to healthcare which enables collaborative model building without sharing sensitive patient data [**fedavg**; Brisimi et al., 2018].

1.3.1 Meta-analysis and Mega-analysis

Meta-analysis and mega-analysis are two well-established techniques for combining and synthesizing information from multiple independent studies. Both methods aim to increase statistical power, improve effect size estimates, and provide more generalizable

results by pooling data from individual studies [Hedges, 1992; Higgins et al., 2019]. However, they differ in their respective approaches and underlying assumptions.

Meta-analysis is a statistical technique that combines the effect size estimates from multiple studies to derive an overall effect size estimate. The method generally involves two steps. First, an effect size measure (e.g., odds ratio, risk ratio, or standardized mean difference) is calculated for each study included in the analysis. Then, a weighted average of these effect size measures is computed, with weights typically assigned based on the inverse of the variance of each study's effect size estimate. This process can be represented mathematically as:

$$\hat{\theta}_{meta} = \frac{\sum_{i=1}^k w_i \hat{\theta}_i}{\sum_{i=1}^k w_i}, \quad (1.1)$$

where $\hat{\theta}_{meta}$ is the overall effect size estimate, k is the number of studies, $\hat{\theta}_i$ is the effect size estimate for study i , and w_i is the weight assigned to study i [Hedges, 1992]. Meta-analysis can employ either fixed-effects or random-effects models, depending on the assumptions made about the true effect sizes and the between-study heterogeneity [Higgins et al., 2019].

Mega-analysis, on the other hand, is a large-scale, collaborative research approach that involves the analysis of individual participant data (IPD) from multiple studies. This approach is particularly useful when the original studies use different outcome measures, as the individual-level data can be harmonized and analyzed in a consistent manner. Unlike meta-analysis, which involves compiling results from published studies, mega-analysis focuses on directly analyzing IPD using an agreed-upon processing strategy. This approach offers several advantages over traditional meta-analysis, such as improved consistency in inclusion criteria across sites, better treatment of confounds and missing data, verification of statistical model assumptions, standardized procedures, increased statistical power, and reduced biases.

Both meta and mega-analysis require the transfer of data to coordinating facilities. In the case of aggregated data, such as histograms of quality metrics, effect sizes, confidence intervals, and standard errors, the risk of reidentification is minimal, as the data is not identifiable at the individual level. However, it is important to note that without proper precautions, repeated computation of aggregate results using slightly varying subsets of participants could expose information about individuals [Dwork, 2006]. This risk can be minimized by establishing agreements among researchers regarding the nature and amount of aggregated data to be transferred.

For mega-analyses, where individual participant data (IPD) is transferred, further attention is required due to differences in regulations across sites that protect the confidentiality, integrity, and security of the IPD and their use in human research. International collaborations like ENIGMA need to accommodate these requirements by adhering to the strictest regulations. Three key aspects for ENIGMA projects include: (a) protection of data and privacy of research subjects, (b) data reuse, and (c) international transfers of data [Zugman et al., 2022].

Mega-analyses with IPD have been shown to provide superior results compared to meta-analyses in terms of higher statistical power and acceptable false-positive rates. However, the major challenge of mega-analysis is the need for at least one site to possess the necessary resources and expertise to handle large datasets. Additionally, this approach is only possible when IPD can be shared with a central facility, which is often limited due to varying data protection regulations among research projects, consortia, and countries [Mathew et al., 1999; Eisenhauer, 2021].

Figure 1.1 illustrates the differences between meta-analyses and mega-analyses in the particular case of the ENIGMA consortium. The ENIGMA (Enhancing Neuro Imaging Genetics through Meta-Analysis) consortium is a global network of researchers that aims to identify genetic and environmental factors that affect brain structure and function using imaging and other measures. The consortium was established in 2009 and has since grown to include over 70 research groups worldwide.

The ENIGMA consortium has produced many influential studies and meta-analyses that have shed light on the genetic and environmental factors that underlie brain development, aging, and disease. For example, a meta-analysis of 94 studies conducted by ENIGMA found that the volume of the hippocampus, a brain region associated with memory and learning, is reduced in individuals with major depressive disorder [Schmaal et al., 2016]. ENIGMA has also investigated the genetic architecture of brain structure and identified novel genetic variants associated with brain volume [Satizabal et al., 2019].

While both meta-analysis and mega-analysis are valuable techniques for synthesizing information from multiple studies, they each have their limitations. Meta-analysis relies on summary statistics from individual studies, which may not capture the full complexity and heterogeneity of the underlying data. Mega-analysis, in contrast, requires access to individual-level data, which can be challenging to obtain and share due to privacy concerns and data protection regulations, such as GDPR and CCPA [Voigt et al., 2017; Solove et al., 2020]. As a result, collaborative efforts like the ENIGMA consortium have demonstrated the potential of multi-centric studies in neuroimaging and started growing interest in exploring alternative approaches, such as federated learning, which enable the development of models compliant with existing regulations without compromising statistical power, generalization, or the need to share raw individual-level data.




Fig. 1.1: Comparison between classical literature-based meta-analyses conducted without access to individual participant data (IPD) (upper panel) and approaches used by different ENIGMA working groups with access to IPD (lower panel). The lower panel shows three main approaches: (top) data processed using common methods at each site, summary statistics computed and sent to a coordinating facility for meta-analysis; (middle) data processed using common methods at each site, sent to the coordinating facility for mega-analysis; and (bottom) raw data sent to the coordinating facility for batch processing and mega-analysis while accounting for site-specific effects. **Reprinted with permission from Andre Zugman, Mega-analysis methods in ENIGMA: The experience of the generalized anxiety disorder working group, 2022.**

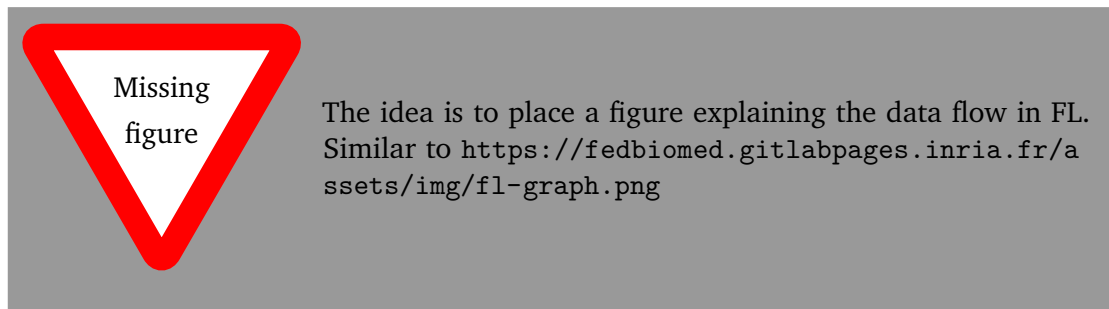


Fig. 1.2: Diagram illustrating the structure of the federated learning framework. The framework adopts a distributed approach to train the model, where participating institutions periodically transmit their locally-updated models to a central server, which in return updates a global model. Such updates are computed by the server (model aggregation) and shares back the aggregated parameters of the global model back to the participating institutions.

1.3.2 Federated Learning

Federated learning (FL) emerges as a recent paradigm that shows promising potential for multi-centric studies in healthcare, as it addresses the data governance and privacy concerns while allowing researchers to collaboratively train models on distributed data.

Federated learning enables multiple parties to collaboratively train a machine learning model without sharing sensitive data [McMahan et al., 2017]. Instead of centralizing data or querying it from all participants, FL performs decentralized optimization where only parameters of a model are allowed to be shared without compromising the privacy of sensitive information. In FL, the model training process typically involves the following steps: [Konečný et al., 2016]:

1. Initialization: A central server initializes the global model parameters and shares them with all participating clients (e.g., healthcare institutions).
2. Local training: Each client trains the model on their local dataset using the current global model parameters, producing local model updates.
3. Aggregation: Clients send their local model updates to the central server, which aggregates these updates to produce a new global model.
4. Global update: The central server shares the updated global model with all clients, and the process repeats from step 2 until convergence. Each repetition of this cycle is considered as a *round* of communication.

As shown in Figure 1.2, two main actors are to be identified in a FL setup: the *clients* and the *server*. Clients, also known as data owners or nodes, are individual entities, such as

healthcare institutions, that possess their local datasets. These clients are connected to a network and participate in the federated learning process by training their local models and communicating with the central server. The central server, sometimes referred to as a third-party aggregator or coordinator, facilitates the federated learning process by initializing the global model, aggregating local model updates from clients, and sharing the updated global model back with the clients. The central server plays a crucial role in managing the communication and coordination among the clients, ensuring that the federated learning process converges to a global model that benefits from the collective knowledge of all participating clients.

Mathematically, federated learning can be formulated as follows. Let \mathcal{D}_k denote the local dataset of the k -th participant, and $f_k(\theta)$ denote the local objective function that measures the quality of the model parameters θ on the local dataset \mathcal{D}_k (e.g. loss function). The global objective function is the average of local objective functions:

$$F(\theta) = \frac{1}{K} \sum_{k=1}^K f_k(\theta), \quad (1.2)$$

where K is the total number of participants. In FL, the goal is to minimize $F(\theta)$ by updating θ through a collaborative optimization process.

A commonly used approach for aggregation federated learning is Federated Averaging (FedAvg) [McMahan et al., 2017], which is a variant of the stochastic gradient descent (SGD) algorithm adapted for the federated setting. In each round of FedAvg, participants compute gradients on their local datasets and send the gradients or model updates to the central server. The server aggregates the updates and computes the global model as follows:

$$\theta^{(t+1)} = \sum_{k=1}^K \frac{|\mathcal{D}_k|}{|\mathcal{D}|} \Delta \theta_k^{(t)}, \quad (1.3)$$

where $\theta^{(t)}$ is the global model at round t , $\Delta \theta_k^{(t)}$ is the update from the k -th participant, and \mathcal{D} is the union of all local datasets where $|\mathcal{D}_k|$ defines the number of observations for participant k .

Despite FedAvg's guarantees and practicality, there have been additional efforts to tackle some of the challenges discussed in Section 1.4, which also depend on the particular use case. Some of these aggregation schemes simply extend the advantages of existing centralized optimizers, such as FedAdam, FedYogi, and FedAdagrad [Reddi et al., 2020],

which extend Adam [Kingma et al., 2014], Yogi [Zaheer et al., 2018], and Adagrad [Lydia et al., 2019], respectively. Other methods, like FedProx [Li et al., 2020a], focus on mitigating the effects of heterogeneous data on the generalizability of the models. Finally, some aggregation schemes aim to reduce communication rounds by scheduling learning rates under different strategies [Chang et al., 2018; Sheller et al., 2019].

1.4 Challenges in Federated Learning and impact on Applications to Healthcare

Despite the considerable potential of federated learning in advancing multi-centric studies in healthcare, FL also faces a unique set of challenges. These challenges stem from the need to maintain data privacy and security, the inherent heterogeneity of healthcare data, communication constraints, and the demands of model convergence and interpretability. Addressing these challenges is critical for the successful application of federated learning to healthcare and the development of reliable, generalizable, and privacy-preserving models.

1.4.1 Framework and Infrastructure

Developing a robust and scalable federated learning framework for healthcare applications is a significant challenge. The framework should accommodate a wide range of data types, models, and optimization methods while ensuring that data privacy and security are maintained. Additionally, the infrastructure should be adaptable to different computational resources, hardware configurations, and network conditions, allowing for seamless collaboration between institutions with varying technical capabilities [Brisimi et al., 2018; Chang et al., 2018].

1.4.2 Communication Efficiency and Resource Utilization

In federated learning, communication between the central server and participating clients (e.g., healthcare institutions) is crucial for model updates and coordination. The communication efficiency directly impacts the overall performance of the federated learning process, particularly in settings with limited bandwidth or unreliable network connections. It is essential to develop communication-efficient federated learning algorithms and techniques that minimize the amount of transmitted information while maintaining model accuracy [Konečný et al., 2016; Li et al., 2019c].

Resource utilization is another critical aspect of federated learning. The computational resources of participating clients can be heterogeneous, ranging from high-performance computing clusters to resource-constrained devices. Efficient resource utilization strategies should be devised to balance the workload and ensure that clients can participate in the federated learning process without overburdening their computational resources [Smith et al., 2017; Lim et al., 2020].

1.4.3 Heterogeneity

The inherent heterogeneity among participating institutions in federated learning can result in biased model updates and suboptimal learning performance, limiting the accuracy and generalizability of the resulting models. Such heterogeneity can come from different data distributions, demographic compositions, and data collection protocols. To address this issue, two main approaches can be taken: developing robust models that can generalize well despite the heterogeneity [Li et al., 2020a; Zhao et al., 2018], or, as is typically suggested in medical studies, correcting biases via harmonization methods [Orlhac et al., 2022].

Harmonization techniques aim to account for variations in the data that can result from differences in data collection, preprocessing, and other factors that could introduce biases. By adjusting the data to minimize these biases, harmonization methods can improve the performance of models trained on multi-center data. This approach is particularly relevant in medical imaging studies, where differences in imaging protocols and scanner hardware can lead to significant variations in the data.

Despite the widespread use of harmonization techniques in medical studies, their adoption in federated learning has been relatively limited. This is because the main objective of harmonization is to remove biases that could lead to misleading findings while preserving the effect of interest. In contrast, machine learning approaches may prioritize better classification or regression performance on unseen data. Currently, the proposals for harmonization remain valid mostly for multi-centric studies centralizing the data. These proposals range from the widely adopted ComBat model in multi-centric medical studies [Johnson et al., 2007], to more complex solutions such as nonlinear versions of it such as ComBat-GAM [Pomponio et al., 2020]. In addition, differently driven approaches like domain adaptation integration using generative adversarial networks (GANs) have also been proposed, as suggested by Wachinger et al. [Wachinger et al., 2021]. Incorporating harmonization techniques in federated learning is a promising area of research that could improve the performance and generalizability of models trained on heterogeneous data.

1.4.4 Privacy and Security

While federated learning is inherently designed to preserve privacy, ensuring the confidentiality and security of sensitive patient data remains a challenge. Potential risks include model inversion attacks, membership inference attacks, and adversarial machine learning attacks. Thus, it is crucial to integrate advanced privacy-preserving mechanisms, such as differential privacy and secure multi-party computation, into the federated learning framework to protect against potential threats [Geyer et al., 2017; Bonawitz et al., 2017].

1.4.5 Model Convergence and Stability

The distributed nature of federated learning can make model convergence and stability more challenging compared to traditional centralized machine learning approaches. Communication delays, resource constraints, and data heterogeneity can lead to slow or inconsistent model convergence. Developing techniques to monitor and improve convergence rates, as well as ensuring model stability under varying conditions, is essential for the practical application of federated learning in healthcare [Li et al., 2018; Sattler et al., 2019].

1.5 Aims and Structure of the Thesis

The primary aim of this thesis is to address the challenges and explore the potential of federated learning in healthcare. We will investigate the development of a robust federated learning framework, optimize resources required for effective federated learning, and improve data harmonization techniques within a federated learning framework, particularly in the context of medical imaging. This work will enable researchers to access and analyze previously inaccessible data in a more robust and generalizable manner while adhering to privacy and ethical standards.

The thesis is structured as follows:

- Chapter 2 introduces Fed-ComBat, a federated approach for harmonization on decentralized data that preserves nonlinear covariate effects.
- In Chapter 3, we present another effort in bias modeling based on our proposed federated mixed-effects variational autoencoders.

- Chapter 4 focuses on optimizing the utilization of communication and computation resources in federated learning through black-box optimization.
- Chapter 5 highlights the need for a common framework for federated learning in healthcare and proposes Fed-BioMed, an open-source infrastructure for federated learning that fulfilling such need.

By addressing these key aspects, this thesis will contribute to the ongoing development and adoption of federated learning in healthcare, enabling more accurate and personalized medicine.

1.6 Publications

As a result of the previously described work, following publications have been achieved:

- **(In submission)** Silva, Santiago, et al. "Fed-ComBat: A Generalized Federated Framework for Batch Effect Harmonization in Collaborative Studies." *Neuroimage*, 2023.
- Silva, Santiago, et al. "Fed-BioMed: A general open-source frontend framework for federated learning in healthcare." *Domain Adaptation and Representation Transfer, and Distributed and Collaborative Learning: Second MICCAI Workshop, DART 2020, and First MICCAI Workshop, DCL 2020, Held in Conjunction with MICCAI 2020, Lima, Peru, October 4–8, 2020, Proceedings 2*. Springer International Publishing, 2020.

1.6.1 Collaborations

During the conception and course of this thesis, various collaborations were pursued, resulting in the following publications:

- Terrail, J. O. D., Ayed, S. S., Cyffers, E., Grimberg, F., He, C., Loeb, R., ... & Andreux, M. (2022). FLamby: Datasets and Benchmarks for Cross-Silo Federated Learning in Realistic Healthcare Settings. *NeurIPS*, 2022.
CRedit Statement: Data Curation, Formal Analysis, Software, Writing - Original Draft.
- Balelli, I., Silva, S., Lorenzi, M., & Alzheimer's Disease Neuroimaging Initiative. (2021). A probabilistic framework for modeling the variability across federated

datasets. In Information Processing in Medical Imaging: 27th International Conference, IPMI 2021, Virtual Event, June 28–June 30, 2021, Proceedings 27 (pp. 701-714). Springer International Publishing.

CRedit Statement: Data curation, Software, Writing - Original Draft.

- S. Silva, B. A. Gutman, E. Romero, P. M. Thompson, A. Altmann and M. Lorenzi, "Federated Learning in Distributed Medical Databases: Meta-Analysis of Large-Scale Subcortical Brain Data," 2019 IEEE 16th International Symposium on Biomedical Imaging (ISBI 2019), Venice, Italy, 2019, pp. 270-274.

CRedit Statement: Conceptualization, Software, Validation, Formal Analysis, Investigation, Writing - Original Draft, Visualization.

FedComBat: Secure Data Harmonization via Federated Learning

2.1	Introduction	22
2.2	Methods	25
2.2.1	Generalized ComBat model	25
2.2.2	Federated ComBat: Fed-ComBat	26
2.2.3	Related works	27
2.3	Materials	30
2.3.1	Synthetic data	30
2.3.2	Brain MRI-Data	32
2.4	Results	34
2.4.1	Synthetic data	34
2.4.2	Brain MRI-data	37
2.5	Limitations and Future Directions	37
2.6	Discussion and Conclusion	41
2.7	Funding	41
2.8	Acknowledgments	42
2.9	Comparing centralized and Fed-ComBat's formulation	43
2.10	Harmonization on synthetic data	43
2.11	Harmonization on real data	44
2.11.1	Evidence of non linear effects on brain phenotypes	44
2.11.2	Residual ComBat effects on real data.	44
2.12	Identifiability of ComBat parameters	47
2.13	CRediT author statement	47

Abstract: In neuroimaging research, the utilization of multi-centric analyses is crucial for obtaining sufficient sample sizes and representative clinical populations. Data harmonization techniques are typically part of the pipeline in multi-centric studies to address systematic biases and ensure the comparability of the data. However, most multi-centric studies require centralized data, which may result in exposing

individual patient information. This poses a significant challenge in data governance, leading to the implementation of regulations such as the GDPR and the CCPA, which attempt to address these concerns but also hinder data access for researchers. Federated learning offers a privacy-preserving alternative approach in machine learning, enabling models to be collaboratively trained on decentralized data without the need for data centralization or sharing.

In this paper, we present Fed-ComBat, a federated framework for batch effect harmonization on decentralized data. Fed-ComBat extends existing centralized linear methods, such as ComBat and distributed as d-ComBat, and nonlinear approaches like ComBat-GAM in accounting for potentially nonlinear and multivariate covariate effects. By doing so, Fed-ComBat enables the preservation of nonlinear covariate effects without requiring centralization of data and without prior knowledge of which variables should be considered nonlinear or their interactions, differentiating it from ComBat-GAM. We assessed Fed-ComBat and existing approaches on simulated data and multiple cohorts comprising healthy controls (CN) and subjects with various disorders such as Parkinson's disease (PD), Alzheimer's disease (AD), and autism spectrum disorder (ASD).

Results indicate that Fed-ComBat outperforms centralized ComBat in the presence of nonlinear effects and is comparable to centralized methods such as ComBat-GAM. Using synthetic data, Fed-ComBat is able to better reconstruct the target unbiased function by 35% (RMSE = 0.5952) with respect to d-ComBat (RMSE = 0.9162) and 12% with respect to our proposal to federate ComBat-GAM, d-ComBat-GAM (RMSE= 0.6751) and exhibits comparable results on MRI-derived phenotypes to centralized methods as ComBat-GAM without the need of prior knowledge on potential nonlinearities.

2.1 Introduction

With the vast generation of neuroimaging data across multiple institutions, concerns seeking to protect sensitive data have been issued as the General Data Protection Regulation - GDPR and the California Consumer Privacy Act - CCPA [European Commission, 2016; Bukaty, 2019; Calzada, 2022]. Such policies demand changes on the paradigms for multicentric clinical studies which nowadays mostly rely on data centralization.

Additionally, data harmonization is known to be a crucial factor in tackling certain challenges of multicentric data analysis due to systematic errors or *batch effects* commonly present in such studies [Wachinger et al., 2021]. Harmonization aims to address biases and variations among data collected from different sources (e.g. scanner brands or

institutions), ensuring that the data is comparable and can be combined effectively for analysis. [Johnson et al., 2007] proposed ComBat in the field of gene expression analysis, as a way to correct these batch effects while preserving the desired covariate effects (e.g., sex, diagnosis, age). Later, this method was adapted to brain magnetic resonance imaging (MRI) phenotypes [Fortin et al., 2017; Fortin et al., 2018]. However, the typical use of ComBat in multicentric studies still requires data to be centralized, thus limiting the application to real-life collaborative analysis scenarios. To overcome this problem, [Chen et al., 2022] proposed d-ComBat to harmonize data within a distributed setup by extending the optimization of ComBat to allow bias correction without the need to exchange the data across centers. To achieve this, d-ComBat shares full local covariance matrices computed at each site, and the subsequently aggregates these covariance matrices to estimate the effect parameters at a global level. This method was shown to achieve almost an exact solution with respect to its centralized version requiring only a few rounds of communication. However, sharing full covariance matrices represents a potential risk of sensitive data leakage [Kesteren et al., 2019], and like ComBat, d-ComBat considers that covariates have a linear influence on phenotypes, which may not always be the case as reported by [Bethlehem et al., 2022] where age shows a clear nonlinear influence on MRI-derived phenotypes. Afterwards, [Pomponio et al., 2020] introduced ComBat-GAM. A method that relies on generalized additive models (GAMs) to address the issue of preserving nonlinear covariate effects. ComBat-GAM uses smoothing functions, such as polynomial or splines, to decompose covariate effects into smooth basis functions, thereby approximating such nonlinear influences. Although ComBat-GAM demonstrates better performance than ComBat, it has limitations such as being designed only for centralized settings. Moreover, ComBat-GAM requires an explicit definition of covariate interactions, typically limited to additive or multiplicative terms (e.g., Sex \times Diagnosis). Despite efforts, a gap still exists in the availability of methods for nonlinear covariate effect preservation without requiring data centralization.

Federated learning (FL) is a popular machine learning paradigm to securely optimize models on decentralized data while preserving institutions' data governance and individuals' sensitive data [Konečný et al., 2016]. Rather than requiring data to be centralized, FL allows institutions to collaboratively train a model keeping data on-site, and sharing only model updates, aggregating these updates on a central server, and distributing the updated global model back to each participant. This iterative process continues until convergence is achieved. Federated Averaging (FedAvg) is a popular approach for aggregation in federated learning. It adapts the stochastic gradient descent (SGD) algorithm for the federated setting providing convergence guarantees [McMahan et al., 2017]. FL has facilitated access to larger cohorts and it has successfully been applied in neuroimaging for tumor segmentation and diagnosis [Li et al., 2019b; Mahlool et al., 2022], disease studies in functional magnetic resonance imaging (fMRI) [Li et al., 2020b], intracranial hemorrhage (ICH) detection [Cheung et al., 2023] and continues to be increasingly adopted in neuroimaging applications.

Within the context of federated learning (FL) and data harmonization, an opportunity opens for the development of flexible and effective federated harmonization approaches that can be easily integrated into FL pipelines, enabling harmonization of heterogeneous data across sites while accounting for complex covariate effects on heterogeneously distributed data. Such approaches have the potential to fill the current gap and allow for more comprehensive analysis of distributed cohorts using FL. To tackle the mentioned challenges, our contributions are considered as follows: i) we introduce a formulation of ComBat which generalizes covariate effect modeling allowing the introduction of more complex functions and, ii) we develop Fed-ComBat, a generalized federated ComBat framework, that allows data batch effect harmonization in federated settings.

We benchmarked and compared our proposed methods with existing centralized approaches (ComBat, and ComBat-GAM) and distributed (d-ComBat) [Johnson et al., 2007; Pomponio et al., 2020; Chen et al., 2022]. Moreover, for the sake of comparison we also formulated and implemented a distributed version of ComBat-GAM (d-ComBat-GAM), which relies on distributed covariance estimation, and extend the method of [Chen et al., 2022] by allowing basis decomposition to interpolate nonlinear covariate effects using smoothing functions.

The different methods were compared for their ability to harmonize batch effects and preserve the quality of covariate effects on simulated data (Section 2.4). Furthermore, we performed an evaluation on derived phenotypes from MRI-brain images from nine cohorts (A4, ABIDE-I, ABIDE-II, ADNI, AIBL, MIRIAD, OASIS3, PDPB, PPMI) comprising controls and people with different brain disorders: patients with different subtypes of Parkinson’s disease (PD), Alzheimer’s disease (AD) and Autism spectrum disorder (ASD).

Using synthetic data, we find that Fed-ComBat achieves a 35% improvement in reconstructing the target unbiased function compared to d-ComBat ($RMSE = 0.9162$), and a 12% improvement compared to our federated extension of ComBat-GAM, d-ComBat-GAM ($RMSE = 0.6751$). In addition, our application on real data shows that Fed-ComBat on MRI-derived phenotypes is comparable to centralized ComBat-GAM. This implies that the studied cohorts may not have exhibited complex nonlinear effects or interactions that could not be captured by a GAM. However, Fed-ComBat requires no prior assumptions of nonlinearities unlike ComBat-GAM and automatically captures such effects and interactions through the model. Moreover, Fed-ComBat can be considered a better approach towards privacy preserving methods as it only shares model parameters, while d-ComBat and our included extension d-ComBat-GAM share full covariance and cross-covariance matrices [Kesteren et al., 2019].

2.2 Methods

2.2.1 Generalized ComBat model

Following the original ComBat formulation proposed by [Johnson et al., 2007], let us denote a batch (represented by different scanners protocols, machines, or institutions) indexed by $i \in \{1, 2, \dots, S\}$ on a particular phenotype (e.g., a brain region) indexed by $g \in \{1, 2, \dots, G\}$. Each batch contains n_i number of observations, and the total number of observations is $N = \sum_i^S n_i$. S can denote for simplicity the number of sites in the study, but it can also be extended to the total number of scanners between sites or any other number of batch effects. We can model a specific phenotype g observed in the j -th patient who belongs to the i -th site denoted by y_{ijg} as follows:

$$y_{ijg} = \alpha_g + \phi(\mathbf{x}_{ij}, \boldsymbol{\theta}_g) + \gamma_{ig} + \delta_{ig}\varepsilon_{ijg}, \quad (2.1)$$

where \mathbf{x}_{ij} denotes the covariate effects expected to be preserved after removing the batch effects (e.g., sex and age), α_g acts as a global fixed intercept (i.e., the mean), while γ_{ig} indicates a random intercept that accounts for the site-specific shift. ε_{ijg} is a noise model that captures the variability of each phenotype $\varepsilon_{ijg} \sim \mathcal{N}(0, \sigma_g^2)$, and δ_{ig} is a multiplicative effect that scales the “unbiased” phenotype variability to fit the one at each site.

This formulation generalizes the original linear model proposed by [Johnson et al., 2007] and the nonlinear covariate effect model proposed by [Pomponio et al., 2020]. While the original model considered $\phi(\mathbf{x}; \boldsymbol{\theta}_g)$ to be a linear function of \mathbf{x} , and the nonlinear model considered $\phi(\mathbf{x}; \boldsymbol{\theta}_g)$ to be an univariate spline, this formulation treats $\phi(\mathbf{x}; \boldsymbol{\theta}_g)$ as a general function that can potentially be multivariable and nonlinear and parametrized by $\boldsymbol{\theta}_g$. This generalization allows for the use of more complex approximating functions such as kernel-based models or neural networks as the multilayer perceptron (MLP), which is known to be a universal function approximator. However, it is essential to use cross-validation techniques to avoid overfitting and ensure the model generalizes well to unseen data.

Some constraints are defined during the fixed effect parameter estimation ($\hat{\alpha}_g, \hat{\boldsymbol{\theta}}_g, \hat{\gamma}_{ig}, \hat{\sigma}_g$):

$$\arg \max_{\hat{\alpha}_g, \hat{\boldsymbol{\theta}}_g, \hat{\gamma}_{ig}, \hat{\sigma}_g} P(y_{ijg} | \hat{\alpha}_g, \hat{\boldsymbol{\theta}}_g, \hat{\gamma}_{ig}, \hat{\sigma}_g) \quad (2.2)$$

$$\text{subject to } \mathbb{E}_g[\hat{\gamma}_i] = \sum_i^S \frac{n_i}{N} \hat{\gamma}_{ig} = 0, \forall g \in \{1, \dots, G\} \quad (2.3)$$

$$\text{and } \phi(\mathbf{x}, \boldsymbol{\theta}_g)|_{x=0} = 0 \quad (2.4)$$

A first constraint in Equation (2.3) is set to allow identifiability of the intercept parameters as explained by [Johnson et al., 2007] (explained in practice in Section 2.12). However, in the same regard, the constraint in Equation (2.4) ensures that no batch effects are captured by the covariate function $\phi(\cdot)$. This second constraint, despite not being mentioned, is also fulfilled by [Johnson et al., 2007] and [Pomponio et al., 2020]; making ComBat and ComBat-GAM a particular case of the proposed formulation in this work.

For a centralized setup, the estimation of all these parameters is performed in three steps: i) maximum likelihood estimation (MLE) for parameters $\hat{\alpha}_g, \hat{\theta}_g, \hat{\gamma}_{ig}$ (see Equation (2.2)) and of the phenotype variance $\hat{\sigma}_g^2 = \frac{1}{N} \sum_{ij} (y_{ijg} - \hat{\alpha}_g - \phi(\mathbf{x}_{ij}; \hat{\theta}_g) - \hat{\gamma}_{ig})^2$, ii) residual standardization mapping the residuals to satisfy the form $y_{ijg} \rightarrow z_{ijg} \sim \mathcal{N}(\gamma_{ig}, \delta_{ig}^2)$ as follows:

$$z_{ijg} = \frac{y_{ijg} - \hat{\alpha}_g - \phi(\mathbf{x}_{ij}; \hat{\theta}_g)}{\hat{\sigma}_g} \quad (2.5)$$

and iii) estimation of the additive and multiplicative batch effects $\hat{\gamma}_{ig}^*$ and $\hat{\delta}_{ig}^*$ as in Equation (2.6), using empirical Bayes (EB) with priors on γ_{ig} and δ_{ig}^2 to iteratively estimate these parameters as in Equation (2.7) [Johnson et al., 2007, Sec. 3.2].

$$\gamma_{ig} \sim \mathcal{N}(Y_i, \tau_i^2) \quad \text{and} \quad \delta_{ig}^2 \sim \text{Inverse Gamma}(\lambda_i, \vartheta_i) \quad (2.6)$$

$$\gamma_{ig}^* = \frac{n_i \bar{\tau}_i^2 \hat{\gamma}_{ig} + \delta_{ig}^{2*} \bar{\gamma}_{ig}}{n_i \bar{\tau}_i^2 + \delta_{ig}^{2*}}, \quad \delta_{ig}^* = \frac{\bar{\theta}_i + \frac{1}{2} \sum_{j=1}^{n_i} (z_{ijg} - \gamma_{ig}^*)^2}{\frac{n_i}{2} \bar{\lambda}_i - 1} \quad (2.7)$$

Lastly, phenotypes can be harmonized while preserving the covariate effects of interest as follows:

$$y_{ijg}^{\text{ComBat}} = \frac{\hat{\sigma}_g}{\hat{\delta}_{ig}^*} (z_{ijg} - \hat{\gamma}_{ig}^*) + \hat{\alpha}_g + \phi(\mathbf{x}_{ij}; \hat{\theta}_g) \quad (2.8)$$

In the following section, we will discuss how this formulation facilitates the incorporation of harmonization within the federated learning framework.

2.2.2 Federated ComBat: Fed-ComBat

Considering the formulation previously presented in Equation (2.1), the parameters α_g, θ_g and γ_{ig} can be optimized by minimizing an objective function F . As data is now

siloed, it is only possible to have an evaluation of the cost function at each site F_i , thus defining the federated optimization problem as:

$$\arg \min_{\alpha_g, \boldsymbol{\theta}_g, \gamma_{ig}} F(\alpha_g, \boldsymbol{\theta}_g, \gamma_{ig}), \quad (2.9)$$

$$\text{where, } F(\alpha_g, \boldsymbol{\theta}_g, \gamma_{ig}) := \sum_{i=1}^S \frac{n_i}{N} F_i(\alpha_g, \boldsymbol{\theta}_g, \gamma_{ig}) \quad (2.10)$$

We rely on *Federated averaging* (FEDAVG) to tackle this optimization problem [McMahan et al., 2017] by allowing each site to conduct partial optimization using stochastic gradient descent (SGD) locally on their data, followed by an aggregation step that combines the shared parameters of each site. The convergence of the resulting iterative process has been demonstrated for both IID and non-IID data distribution across clients [Li et al., 2019c].

The Fed-ComBat framework requires the following steps for harmonization: i) a federated standardization step to avoid scaling issues in gradient descent as proposed by [Silva et al., 2019], ii) a federated estimation of $\alpha_g, \boldsymbol{\theta}_g$ and γ_{ig} as presented in Equation (2.10), and iii) a local estimation of the random effects following using EB as in Equations (2.6) and (2.7). A description of the steps followed in federated harmonization using Fed-ComBat is described in Algorithm 1 including standardization as preprocessing, local updates relying on SGD, global updates across sites of shared parameters using FEDAVG and random effect estimation using Empirical Bayes.

2.2.3 Related works

d-ComBat

[Chen et al., 2022] proposed a distributed version of ComBat that allows correction for batch effects in a distributed manner. This is achieved by taking advantage of the multivariate linear regression problem solved in ComBat, requiring to compute the product between the inverse of a covariance matrix and the data-targets cross-covariance. Both can be computed by aggregating the individual matrices provided by each site. This allows us to reformulate Equation (2.1) as done by [Chen et al., 2022]:

Algorithm 1: Fed-ComBat

Data: Non-harmonized phenotypes (y_{ijg}) and covariates (\mathbf{x}_{ij}).

Result: Harmonized phenotypes with siloed data.

$\mathbf{x} \leftarrow \text{FEDERATEDSTANDARDIZATION}(\mathbf{x});$

// Estimation of fixed effects and random intercept

initialization of parameter space $\Omega := \{\theta_g, \alpha_g, \gamma_{ig}\}_{g=1}^G;$

while not converged do

foreach site i do

$\Omega_i^t \leftarrow \Omega^{(t+1)};$ // Initialization.

 // Partial local optimization using SGD.

foreach local gradient step t do

$\Omega_i^{(t+1)} = \Omega_i^{(t)} - \eta \nabla_{\Omega_i} F(\Omega_i^{(t)});$

 // Aggregate and update every parameter using FedAVG.

$\Omega^{(t+1)} \leftarrow \sum_i \frac{n_i}{N} \Omega_i^t$

forall site i do

 // Standardize using estimated parameters

$z_{ijg} \leftarrow \frac{y_{ijg} - \hat{\alpha}_g - \phi(\mathbf{x}_{ij}; \hat{\theta}_g)}{\hat{\sigma}_g}$

 // Estimate γ_{ig}^* and δ_{ig}^* using EB

$\gamma_{ig}^*, \delta_{ig}^* \leftarrow \text{EMPIRICALBAYES}(\mathbf{x});$ // Equations (2.6) and (2.7)

 // Correct data

return $y_{ijg}^{\text{ComBat}} \leftarrow \frac{\hat{\sigma}_g}{\delta_{ig}^*} (z_{ijg} - \hat{\gamma}_{ig}^*) + \hat{\alpha}_g + \phi(\mathbf{x}_{ij}; \hat{\theta}_g)$

$$\underbrace{\begin{bmatrix} y_{11g} \\ \vdots \\ y_{ijg} \\ \vdots \\ y_{Sn_s g} \end{bmatrix}}_{\mathbf{y}_g} = \underbrace{\begin{bmatrix} x_{111} & \dots & x_{11C} & | & 1 & \dots & 0 & | & 1 \\ \vdots & & x_{ijc} & | & \vdots & \ddots & \vdots & | & \vdots \\ x_{Sn_s 1} & \dots & x_{Sn_s C} & | & 0 & \dots & 1 & | & 1 \end{bmatrix}}_{\mathbf{X}} \underbrace{\begin{bmatrix} \hat{\theta}_{g1} \\ \vdots \\ \hat{\theta}_{gc} \\ \hat{\gamma}_{1g} \\ \vdots \\ \hat{\alpha}_g \end{bmatrix}}_{\hat{\Theta}_g} + \delta_g \epsilon_g \quad (2.11)$$

Starting from the panel formulation where \mathbf{y}_g is the stacked vector of all subjects across all sites i , and \mathbf{X} is a design matrix containing the covariate effects to be preserved indexed by c in subject j from site i (x_{ijc}), an indicator matrix encoding the site, and a column of ones to capture α_g , we can estimate the augmented parameter matrix $\hat{\Theta}_g$ using maximum likelihood. The estimator can be decomposed as a sum of the covariance and the cross-covariance matrices at each site, as shown in Equation 2.12. This formulation has also been used to distribute latent variable models, such as PCA and PLS, among multiple sites [Silva et al., 2019; Lorenzi et al., 2018]. The authors have also suggested sharing a partial-eigen decomposition to mitigate the risk of data leakage, as shown in Equation 2.13.

$$\hat{\Theta}_g = (\mathbf{X}^\top \mathbf{X})^{-1} \mathbf{X}^\top \mathbf{y}_g = \left(\sum_{i=1}^S \mathbf{X}_i^\top \mathbf{X}_i \right)^{-1} \left(\sum_{i=1}^S \mathbf{X}_i^\top \mathbf{y}_{ig} \right) \quad (2.12)$$

$$= \left(\sum_{i=1}^S \mathbf{Q}_i \Lambda_i \mathbf{Q}_i^{-1} \right) \left(\sum_{i=1}^S \mathbf{U}_i \Sigma_i \mathbf{V}_i^\top \right) \quad (2.13)$$

d-ComBat-GAM

Seeking a federated nonlinear approach, we here propose d-ComBat-GAM as an adaptation for d-ComBat for non-linear covariate effect preservation via generalized additive models (GAM) as proposed in its centralized version by [Pomponio et al., 2020]. This method, as d-ComBat, allows a fast estimation in a closed form by sharing either full or partial decomposition of the covariance and cross covariance matrices but transforming covariates in a design matrix using smooth functions such as polynomials or splines.

We can consider d-ComBat-GAM as a particular case of Fed-ComBat where $\phi(\cdot)$ becomes a linear combination parametrized by θ . Let $b(\cdot)$ be an arbitrary basis representation, the covariate function approximate is then defined as follows:

$$\phi(\mathbf{x}_{ij}; \theta_g) = \sum_{c=1}^C \sum_{k=1}^{K_c} \theta_{gk} b_{gk}(x_{ijc}) \quad \text{where} \quad \theta_g = [\theta_{g1}, \dots, \theta_{gK_c}] \quad (2.14)$$

Note that Equation (2.14) expresses a linear combination of parameters, hence the same procedure as in Section 2.2.3 can be followed to estimate the fixed effects. Although polynomial decomposition is straightforward as it is the equivalent of transforming and appending the variable to a design matrix, splines require a set of defined control points based on the covariate range. We rely on federated standardization to map the nonlinear covariate to a normalized distribution where the range is known and thus the control points of the splines can be predefined.

2.3 Materials


We evaluated Fed-ComBat on synthetic data accounting for different sources of batch- and covariate-wise heterogeneity.

We also benchmarked this approach on a collection of nine cohorts corresponding to different studies in neurodegenerative disorders including participants diagnosed with or at risk of developing Parkinson’s and Alzheimer’s disease, and Autism spectrum disorder as well as controls.

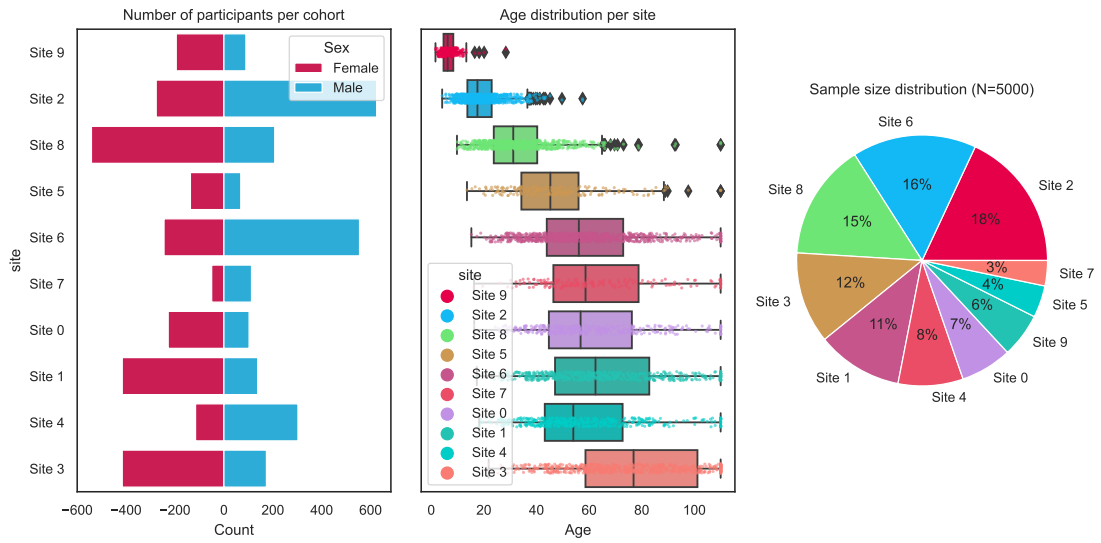
2.3.1 Synthetic data

We created simulated datasets under three main scenarios: i) when cohorts have similar distributions for covariates and sample size effects (IID), ii) when covariate distributions slightly overlap but follow different patterns, and when there are differences in sample sizes (mid-non-IID), and iii) when there is no overlap between covariate distributions between cohorts and there are evident differences in sample sizes (non-IID). We simulated ComBat parameters following the methodology outlined in [Reynolds et al., 2022], using the graphical model depicted in Figure 2.2. The primary difference between our approach and that of [Reynolds et al., 2022] was the model used to relate the covariates and phenotypes ϕ . We used a nonlinear function in this step, which was employed as the unbiased target function to be learned after the harmonization process for the synthetically biased phenotypes.

Sex was generated by sampling from a Bernoulli($p_i^{(\text{sex})}$) distribution, which modulated sex proportion distribution across sites. Age was drawn from a $\mathcal{N}(\mu_i^{(\text{age})}, \sigma_i^{(\text{age})^2})$ distribution, where $\mu_i^{(\text{age})}$ and $\sigma_i^{(\text{age})}$ were sampled from a Uniform(a, b) distribution, with heterogeneity modulated by the width of the distribution ($b - a$). Sample size heterogeneity was simulated using a concentration model $p_i \sim \text{Dir}(\alpha)$, as previously suggested for measuring the effects of heterogeneity in Federated Learning by [Hsu et al., 2019].



(a)



(b)

Fig. 2.1: Population distribution of synthetic data across cohorts in two different scenarios. The left panel shows the sex distribution, the middle panel shows the age distribution, and the right panel shows the sample size distribution. Panel 2.1a shows a homogeneous and IID population, while panel 2.1b shows a non-IID distribution of covariates, which is more commonly observed and used in this work.




Fig. 2.2: Graphical model used to generate synthetic data. The shaded circles indicate observed measurements, including covariates and imaging feature values, while unshaded circles represent latent parameters.

2.3.2 Brain MRI-Data

We evaluated our method on a cross-sectional cohort of 7265 participants from nine public studies, utilizing MRI-derived phenotypes from baseline structural T1-weighted magnetic resonance imaging (MRI).

Participants were categorized by diagnosis as follows: 3992 cognitively normal (CN), of which 731 of them were considered at risk of cognitive impairment from the A4 study [Sperling et al., 2014]; 1065 were identified as having mild cognitive impairment (MCI) from the Alzheimer’s Disease Neuroimaging Initiative (ADNI) [Weiner et al., 2013] and the Australian Imaging, Biomarker and Lifestyle Flagship Study of Aging (AIBL) [Ellis et al., 2009]; 998 had been diagnosed with Autism Spectrum Disorder (ASD) from ABIDE-I and ABIDE II [Di Martino et al., 2014]; and 538 were patients diagnosed with Alzheimer’s disease (AD) from ADNI and Open Access Series of Imaging Studies (OASIS) dataset [Marcus et al., 2010] and, 686 participants diagnosed in multiple phases of Parkinson’s disease from the Parkinson’s Disease Biomarkers Program (PDPB) [Rosenthal et al., 2016] and the Parkinson’s Progression Markers Initiative (PPMI) [Marek et al., 2018]. Further demographic information can be found in Table 5.2.

Image processing

Cortical thickness and subcortical volume were extracted from MRI scans using FreeSurfer v7.1.1 (documented and freely available for download online at: <http://surfer.nmr.mgh.harvard.edu/>) [Dale et al., 1999; Fischl et al., 2004]. Steps for phenotypical measures extraction included skull stripping [Ashburner et al., 2005], Non-uniform intensity Normalization (N3) [Sled et al., 1998], segmentation using the Desikan–Killiany atlas [Desikan et al., 2006] and extraction of the MRI-derived phenotypes.

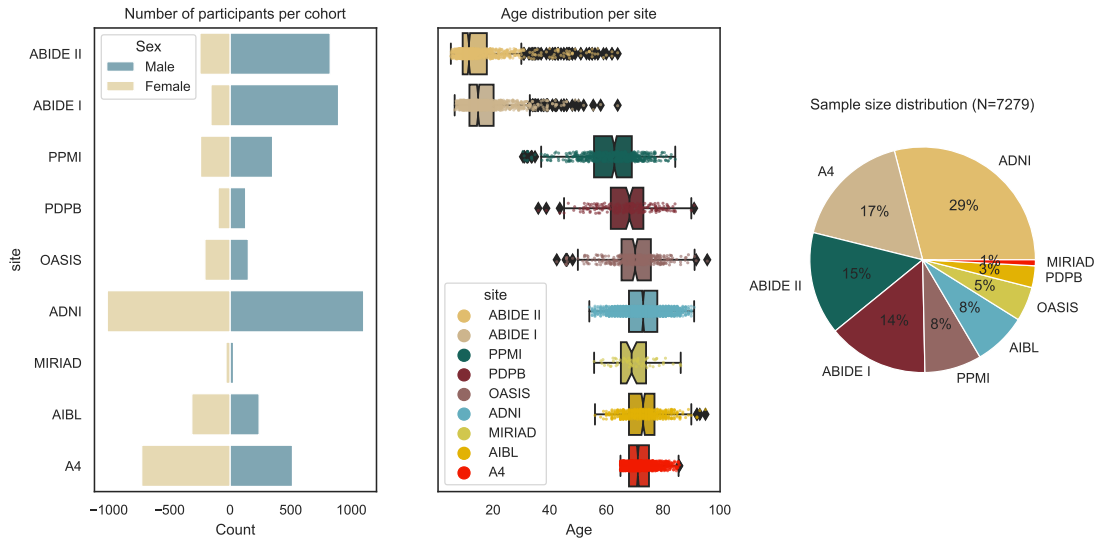


Fig. 2.3: Population pyramid representing the demographics of the real data used in this study. The left panel shows the distribution of sex, the middle panel shows the distribution of age, and the right panel shows the sample size distribution. Cohorts were sorted by ascending median age. The demographics for the population are described in Table 5.2. The figure shows a clear non-IID distribution, similar to the distribution observed in Figure 2.1b.

Harmonization was carried out preserving age, sex, diagnosis and intracranial volume (ICV / eTIV), and each study was considered a site whose systematic biases are expected to be corrected [Reynolds et al., 2022].

site	Group \ Sex	N		Age in years \pm SD [age range]			
		Female	Male	Female		Male	
A4	CN	728	515	71.3 \pm 4.5	[65.0 - 85.7]	72.9 \pm 5.1	[65.0 - 85.7]
ABIDE I	Autism	59	434	16.3 \pm 8.1	[8.1 - 45.0]	17.4 \pm 8.6	[7.0 - 64.0]
	CN	98	446	15.4 \pm 6.6	[7.8 - 46.0]	17.4 \pm 7.6	[6.5 - 48.0]
ABIDE II	Autism	74	416	13.1 \pm 8.1	[5.2 - 54.0]	15.0 \pm 9.3	[5.1 - 62.0]
	CN	175	397	13.5 \pm 7.4	[5.9 - 46.6]	15.6 \pm 10.0	[5.9 - 64.0]
ADNI	AD	156	190	73.6 \pm 7.8	[55.0 - 91.0]	75.0 \pm 7.6	[55.0 - 90.0]
	CN	444	345	71.8 \pm 6.2	[55.0 - 90.0]	74.0 \pm 6.1	[60.0 - 90.0]
	MCI	406	555	71.5 \pm 7.8	[55.0 - 88.0]	73.6 \pm 7.2	[54.0 - 90.0]
AIBL	AD	34	25	75.6 \pm 8.1	[58.0 - 93.0]	72.7 \pm 6.9	[60.0 - 83.0]
	CN	237	169	71.9 \pm 6.0	[60.0 - 89.0]	73.0 \pm 6.3	[60.0 - 90.0]
	MCI	44	46	75.8 \pm 8.4	[56.0 - 95.0]	74.1 \pm 5.6	[64.0 - 87.0]
MIRIAD	AD	22	17	69.5 \pm 6.3	[58.1 - 80.2]	69.4 \pm 7.5	[55.7 - 86.1]
	CN	11	12	66.3 \pm 4.9	[58.8 - 73.8]	73.6 \pm 7.2	[63.5 - 86.3]
OASIS	AD	45	41	75.6 \pm 6.8	[62.9 - 95.6]	75.2 \pm 7.3	[60.5 - 91.7]
	CN	162	109	67.6 \pm 8.6	[45.7 - 88.8]	69.1 \pm 8.3	[42.5 - 86.2]
PDPB	PD	97	129	67.8 \pm 9.1	[50.7 - 90.0]	67.2 \pm 9.7	[36.0 - 91.0]
	CN	51	84	59.1 \pm 11.6	[31.0 - 81.9]	60.8 \pm 11.7	[30.6 - 82.8]
PPMI	GenCohort PD	29	27	65.2 \pm 8.0	[51.7 - 81.2]	63.9 \pm 9.3	[32.2 - 78.6]
	GenCohort Unaff	52	35	61.1 \pm 8.2	[33.7 - 84.3]	61.3 \pm 7.6	[46.8 - 75.0]
	PD	94	164	61.2 \pm 10.2	[33.5 - 81.7]	62.2 \pm 9.4	[34.8 - 82.9]
	Prodromal	2	14	70.1 \pm 4.3	[67.1 - 73.2]	69.0 \pm 6.7	[61.7 - 82.5]
	SWEDD	15	26	61.7 \pm 8.9	[46.8 - 77.6]	61.9 \pm 9.9	[39.2 - 77.0]

Tab. 2.1: Subject demographics of real data used in this work.

2.4 Results

Two versions of Fed-ComBat were used for comparison: a first one defining a linear model ($\phi(\mathbf{x}; \theta) = \mathbf{x}^T \theta$) to show equivalence with respect to ComBat and d-ComBat, and a second one defining ϕ as a multi-layer perceptron model (MLP) for nonlinear effect preservation. These two versions of Fed-ComBat were tested jointly with our extended proposal d-ComBat-GAM against their two centralized versions: ComBat¹ (a.k.a. NeuroComBat), ComBat-GAM² and the only reported distributed version d-ComBat-GAM³.

2.4.1 Synthetic data

Reconstruction of an unbiased phenotype function depending on age and sex after harmonizing using different methods is shown in Figure 2.4a. As observed, centralized ComBat, d-ComBat, and Fed-ComBat in its linear version follow similar trends. It can also be observed that for nonlinear covariate effects linear correction can even induce biases

¹<https://github.com/Jfortin1/neuroCombat>

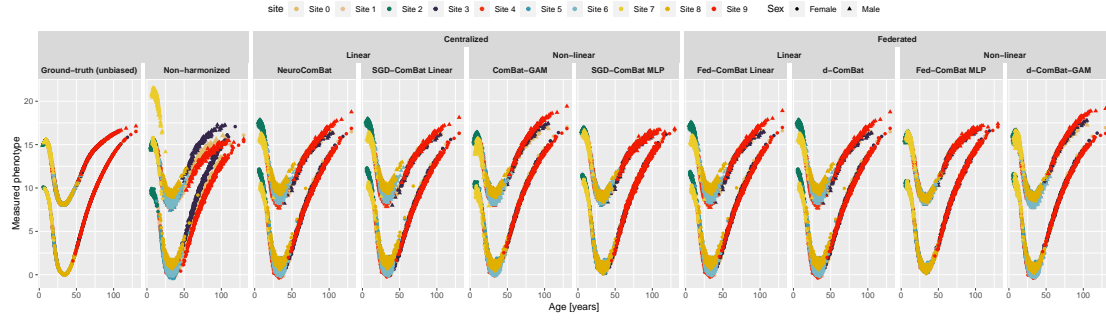
²<https://github.com/rpomponio/neuroHarmonize>

³<https://github.com/andy1764/Distributed-ComBat>


due to its limitation to capture complex relationships between covariates and phenotypes (see Non-harmonized vs. NeuroCombat in Figure 2.4a).

While both ComBat-GAM and our distributed adaptation d-ComBat-GAM demonstrate improved correction for nonlinear covariate effects compared to ComBat (as shown in Figure 2.4a), it is worth noting that d-ComBat-GAM requires a priori definition of covariate interactions. In cases where these interactions are not defined, the performance of d-ComBat-GAM is reduced to that of ComBat. A better reconstruction of the target unbiased trajectories is achieved with Fed-ComBat using a two-layer perceptron neural network with 100 hidden units and with no hyperparameter tuning. Hypertuning was out of the scope of this work. Root mean square errors (RMSE) evaluating the reconstruction quality is presented in Table 2.2 for all centralized and federated methods.

Bland-Altman plots contrasting the difference between the different harmonization methods and the ground truth are presented in Figure 2.4b highlighting that there is less error variability (σ_ϵ) in estimating the unbiased function after harmonization for nonlinear methods and, particularly for Fed-ComBat in terms of the Root mean square error (RMSE = 0.5952) compared to the second federated best d-ComBat-GAM (RMSE = 0.6751).



(a) Reconstruction of an unbiased groundtruth function.



(b) Bland-Altman plots and root mean square errors (RMSE) contrasting the different centralized and federated harmonization methods against the groundtruth.

Fig. 2.4: Qualitative comparison across harmonization methods evaluating the quality of harmonization for a simulated phenotype with a nonlinear relationship with age and following a different trajectory per group (males and females). Results are shown for the extreme non-IID covariate distribution is depicted in Figure 2.1b. This can be seen as certain centers contain extremely young or old cohorts (tails). (a) Expected result after harmonization: groundtruth (unbiased), worst case scenario: biased data (non-harmonized), and harmonized phenotypes using the different ComBat methods considered and proposed in this work. As shown, ComBat could even insert biases when data is too heterogeneous. Fed-ComBat preserves better the real trajectories than the linear approaches and ComBat-GAM without the need of defining where the nonlinearities may be. Fed-ComBat using a multi-layer perceptron (MLP) shows the best reconstruction with an improvement of 35% in RMSE with respect to d-ComBat and 12% with respect to d-ComBat-GAM.

Setup	Approach	Method	$ \bar{\varepsilon} $ (MAE)	σ_{ε}	RMSE
Federated	Non-linear	Fed-ComBat MLP	0.55710	0.20960	0.59520
Federated	Non-linear	d-ComBat-GAM	0.55160	0.38930	0.67510
Centralized	Non-linear	SGD-ComBat MLP	0.55420	0.42150	0.69620
Centralized	Non-linear	ComBat-GAM	0.55530	0.46120	0.72180
Federated	Linear	Fed-ComBat Linear	0.56540	0.70020	0.89990
Centralized	Linear	NeuroComBat	0.55230	0.73110	0.91620
Federated	Linear	d-ComBat	0.55230	0.73110	0.91620
Centralized	Linear	SGD-ComBat Linear	0.58640	0.81000	0.99990

Tab. 2.2: Quantitative results for Figure 2.4a in terms of mean absolute error between methods and the target hidden unbiased trajectories ($|\bar{\varepsilon}|$), its standard deviation (σ_{ε}) and root mean square error (RMSE). Highlighted values correspond to the best metric across different methods.

2.4.2 Brain MRI-data

Seeking evidence of nonlinear covariate effects of age on brain phenotypes, we compared the goodness of fit of two models: a linear model and a GAM. The criterion used to evaluate the presence of nonlinearities was the difference in the Akaike Information Criterion (AIC) between the GAM and the linear model. When the difference is negative, it indicates that the GAM is a better fit, and the magnitude of the difference indicates by how much. Figure 2.5b shows the AIC metric across the regions of the brain using the Desikan-Killiany parcellation and cortical thickness as the endogenous variable and the ASEG atlas for subcortical volumes. The data was controlled for sex, diagnosis, ICV, and group, and age was considered an exogenous variable. The data was also previously harmonized using ComBat. Figure 2.5a shows the top three regions that are better explained by a nonlinear model of age in terms of AIC, as well as the residuals after the regression illustrating remaining effects or trends. More detailed results on the AICs differences is presented in Table 2.3.


Brain trajectories segregated using diagnosis as Normal and Abnormal for ventral diencephalon, right hippocampus and hemisphere rostral middle frontal gyrus before and after harmonization are shown in Figure 2.7. The harmonized phenotypes were obtained using as linear centralized method: ComBat, nonlinear centralized: d-ComBat, linear federated: d-ComBat and our adaptation Fed-ComBat using a linear model, and nonlinear federated: our extension d-ComBat-GAM and Fed-ComBat using an MLP with two layers and 100 hidden units.

The residual harmonized term $\frac{\hat{\sigma}_g}{\hat{\delta}_{ig}^*} (z_{ijg} - \hat{\gamma}_{ig}^*)$ should follow a normal standard noise distribution by definition and should not have center residual effects. Such residuals for the top three regions of interest in Figure 2.5b are illustrated per cohort in Figure 2.6 for the decentralized harmonization methods considered and proposed in this work.


2.5 Limitations and Future Directions

While Fed-ComBat has demonstrated promising results for data harmonization in a federated learning setting, there are potential limitations and challenges that should be considered, as well as opportunities for future research directions.

Model architecture: The choice of model architecture, such as the two-layer perceptron function employed in the current implementation of Fed-ComBat, can influence the performance of the harmonization process. In order to better capture complex nonlinearities in various harmonization problems, it may be useful to explore alternative architec-



- (a) This figure displays the regions of the brain where a generalized additive model (GAM) provides a better fit than a linear model, as measured by the difference in Akaike Information Criterion (AIC) between the two models. The parameters for the GAM are set as in [Pomponio et al., 2020], with age as a smoothing term, B-splines with 10 degrees of freedom or control points uniformly distributed between the minimum and maximum values, and a maximum polynomial of 3rd degree. A negative difference indicates a better fit with the GAM, while a positive difference indicates a better fit with the linear model.



- (b) Results of regression using a GAM on the top 3 regions with the most discrepant AIC values on the left and residual plots illustrating goodness of fit on the right. Top: Ventral diencephalon, middle: Right hippocampus, bottom: Right hemisphere rostral middle frontal gyrus.

Fig. 2.5: Evidence of nonlinear effects present in used cohorts. Subjects were harmonized by cohort using ComBat [Fortin et al., 2018].




Fig. 2.6: Residual terms after harmonization for top three regions discusses in Figure 2.5a for federated methods. Additional comparisons using all the methods are presented in Section 2.11.

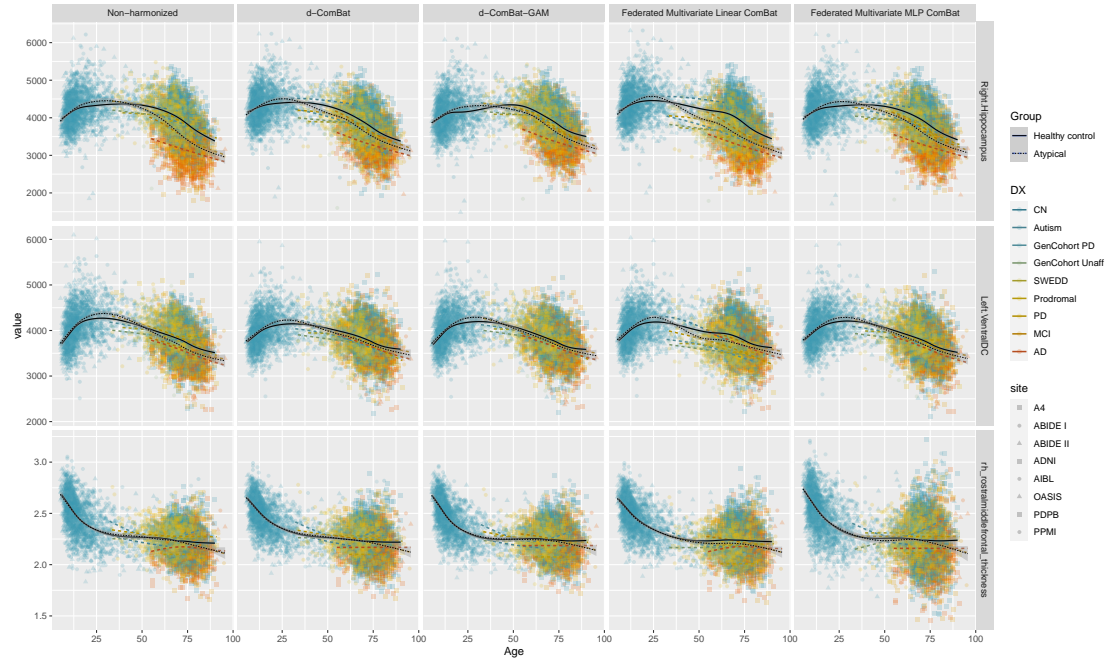


Fig. 2.7: Harmonized phenotypes mentioned in Figure 2.5a in each row. Non-harmonized, d-ComBat and all the proposed methods in this work in each column. Trajectories are drawn after accounting for sex, and ICV after harmonization preserving age, sex, diagnosis and ICV.

tures and conduct hyperparameter tuning if necessary, while respecting the constraints required for integration within the Fed-ComBat framework (see Equation (2.4)).

Heterogeneity: The performance of Fed-ComBat may be affected by the level of heterogeneity across participating institutions. In cases where there is extreme disparity in data distributions, demographic compositions, or data collection protocols, the method might struggle to correct biases effectively. Therefore, the suitability of Fed-ComBat or any ComBat method in such cases should be carefully evaluated. Alternative directions could include exploring optimization methods such as FedProx[Li et al., 2020a].

Computational resources and connectivity: The current implementation of Fed-ComBat relies on the assumption that the participating institutions have enough computational resources and connectivity to perform federated learning. In practice, resource constraints or network issues might limit the applicability of this method, particularly in low-resource settings or in cases where institutions have unreliable or slow network connections.

Application to imaging data: The current implementation of Fed-ComBat focuses on derived phenotypes just like the first version of ComBat [Johnson et al., 2007], but future work could explore the application of this method directly to imaging data. By making $\phi(\mathbf{x}_{ij}, \theta_g)$ a convolutional neural network (CNN), Fed-ComBat could be extended to

handle image data directly, which would open up new opportunities for harmonization in multi-centric imaging studies in federated setups.

By addressing these limitations and challenges, and exploring potential future directions, Fed-ComBat could be further improved and extended to a wider range of applications, ultimately enhancing its utility in federated learning and data harmonization for healthcare research.

2.6 Discussion and Conclusion

In conclusion, we have shown that it is possible to achieve acceptable results compared to the linear case without the need to share full covariance matrices or assume nonlinearities. The proposed Fed-ComBat method offers a novel, generalized approach for data harmonization, inspired by ComBat, and is suitable for scenarios where data sharing is restricted and nonlinear covariate effects might be present. By relying on federated learning, Fed-ComBat can capture these nonlinearities and interactions through its fully connected architecture, allowing for comparable results with centralized approaches and d-ComBat while keeping data decentralized and potentially avoiding data leakage.

In contrast to ComBat-GAM, Fed-ComBat does not require predetermined definitions of variables to be decomposed or possible interactions across variables (i.e., smoothing factor). This is because complex approximation functions like multi-layer perceptrons can effectively capture nonlinearities and interactions without explicit definitions.

As a future direction, it would be beneficial to explore the performance of Fed-ComBat on larger cohorts, as in [Bethlehem et al., 2022], where nonlinear effects are better exacerbated. This would provide further evidence of the method's effectiveness in handling complex harmonization problems in neuroimaging research.

2.7 Funding

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 847579 (Marie Skłodowska-Curie Actions), and by the INRIA Sophia Antipolis-Méditerranée "NEF" computation cluster. This study was partially supported by the Early Detection of Alzheimer's Disease Subtypes (E-DADS) project, an EU Joint Programme - Neurodegenerative Disease Research (JPND) project (see www.jpnd.eu). The project is supported under the aegis of JPND through the following funding organizations: United Kingdom, Medical Research Council (MR/T046422/1); Netherlands, ZonMW (733051106); France, Agence Nationale de

la Recherche (ANR-19-JPW2-000); Italy, Italian Ministry of Health (MoH); Australia, National Health & Medical Research Council (1191535); Hungary, National Research, Development and Innovation Office (2019-2.1.7-ERA-NET-2020-00008).

2.8 Acknowledgments

Data used in preparation of this article were obtained from the Alzheimer's Disease Neuroimaging Initiative (ADNI) database (<https://adni.loni.usc.edu/>). The investigators within the ADNI contributed to the design and implementation of ADNI and/or provided data but did not participate in analysis or writing of this report. A complete listing of ADNI investigators can be found at https://adni.loni.usc.edu/wp-content/uploads/how_to_apply/ADNI_Acknowledgement_List.pdf. Data used in the preparation of this article was obtained from the Australian Imaging Biomarkers and Lifestyle flagship study of ageing (AIBL) funded by the Commonwealth Scientific and Industrial Research Organisation (CSIRO) which was made available at the ADNI database (<https://adni.loni.usc.edu/aibl-australian-imaging-biomarkers-and-lifestyle-study-of-ageing-18-month-data-now-released/>). Data used in the preparation of this article were obtained from the Alzheimer's Disease Neuroimaging Initiative (ADNI) database (<https://adni.loni.usc.edu/>). Data used in this article were provided by the Australian Imaging Biomarkers and Lifestyle flagship study of ageing (AIBL) funded by the Commonwealth Scientific and Industrial Research Organisation (CSIRO) which was made available at the ADNI database (<https://adni.loni.usc.edu/aibl-australian-imaging-biomarkers-and-lifestyle-study-of-ageing-18-month-data-now-released/>). Data used in the preparation of this article were obtained from the Autism Brain Imaging Data Exchange (ABIDE) I database. Data used in this article were obtained from the Autism Brain Imaging Data Exchange (ABIDE) II database (http://fcon_1000.projects.nitrc.org/indi/abide/abide_II.html). The AIBL researchers contributed data but did not participate in the analysis or writing of this report. AIBL researchers are listed at <https://www.aibl.csiro.au>. Data used in the preparation of this article were obtained from the MIRIAD database. Data used in the preparation of this article were obtained from the Open Access Series of Imaging Studies (OASIS) database. Data used in the preparation of this article were obtained from the Parkinson's Progression Markers Initiative (PPMI) database (<http://www.ppmi-info.org/>). Data used in the preparation of this article were obtained from the Pre-symptomatic Evaluation of Experimental or Novel Treatments for Alzheimer's Disease (A4) study.

2.9 Comparing centralized and Fed-ComBat's formulation

As the random effects (batch effects) in this work are assumed to be site-specific, the model's design matrix \mathbf{X} is constructed under a linear assumption, jointly optimizing the parameters $\hat{\alpha}_g$ and $\hat{\gamma}_{ig}$. However, for nonlinear scenarios, the model needs to be reformulated, as previously done by [Pomponio et al., 2020]. To provide a more comprehensive understanding of the optimization problem, we propose the following generalized approach:

$$y_{ijg} = \phi(\mathbf{x}_{ij}; \boldsymbol{\theta}_g) + \mathbf{z}_{ij}\boldsymbol{\beta}_g + \delta_{ig}\varepsilon_{ijg} \quad (2.15)$$

Where with respect to the model in Equation (2.1), $\mathbf{Z} := \{\mathbf{z}_{ij}\}_{\forall i,j}$ is an indicator matrix as indicated in Equation (2.16) on the batch effects that allows the elements of $\boldsymbol{\beta}_g$ be composites of α_g and γ_{ig} as shown in Section 2.9.

$$\mathbf{z}_{ij(1,S)} := \begin{cases} 1 & \text{if } \mathbf{x}_{ij} \in \text{batch } i, \\ 0 & \text{if } \mathbf{x}_{ij} \notin \text{batch } i. \end{cases} \quad (2.16)$$

$$\boldsymbol{\beta}_g = \alpha_g + \begin{pmatrix} \gamma_{1g} \\ \vdots \\ \gamma_{Sg} \end{pmatrix} \quad (2.17)$$

$$\text{s.t. } \mathbb{E}[\gamma_i] = \sum_{i=1}^S \frac{n_i}{N} \gamma_{ig} = 0 \quad \forall g \in \{1, \dots, G\} \quad (2.18)$$

Which with the constraint in Eqn 2.18 the estimation of α_g and γ_{ig} from $\boldsymbol{\beta}_g$ becomes straightforward:

$$\alpha_g = \mathbb{E}[\boldsymbol{\beta}_g] = \sum_{i=0}^S p_i \beta_{ig} = \sum_{i=0}^S \frac{n_i}{N} \alpha_g + \mathbb{E}[\gamma_{ig}] \quad (2.19)$$

$$\gamma_{ig} = \beta_{ig} - \alpha_g \quad (2.20)$$

2.10 Harmonization on synthetic data

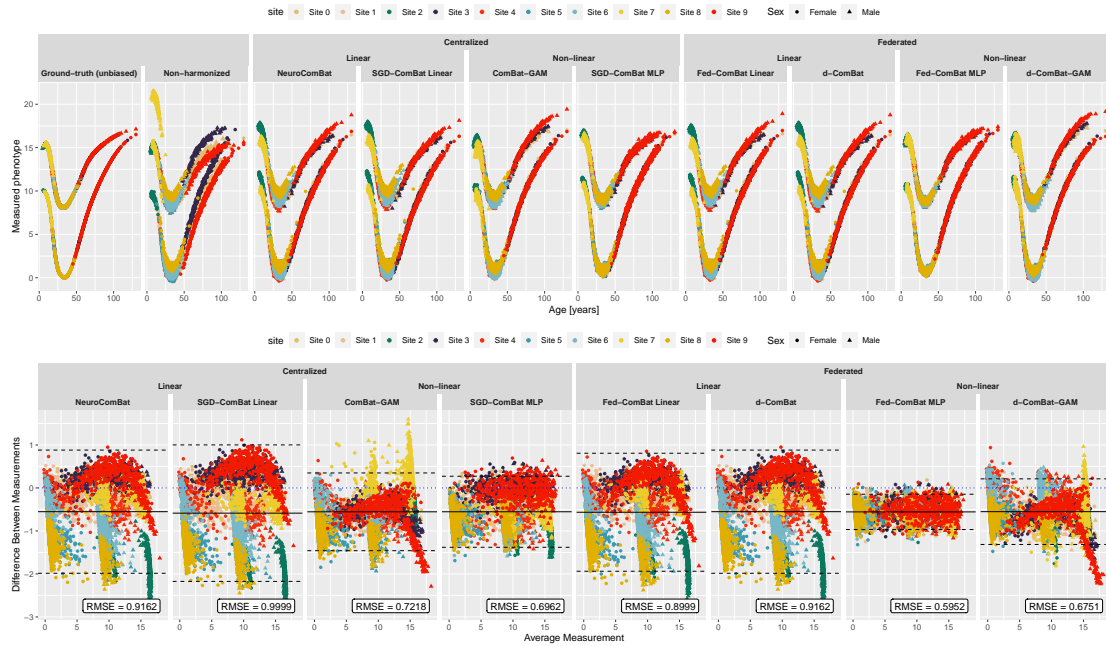


Fig. 2.8: Comparison between centralized and federated methods on synthetic data. The objective is to reconstruct the unbiased (ground-truth) function. As well as their respective bland Altman plots.

2.11 Harmonization on real data

2.11.1 Evidence of non linear effects on brain phenotypes

Table 2.3 shows the difference in Akaike information criterion (AIC) between a GAM and a GLM approximating the phenotypes (brain measures) as proposed by [Pomponio et al., 2020] (B-Splines for approximating age with DOF=10 and degree 3) as a function of Age, Sex, ICV and group (diagnosis). Negative values indicate a better fit using GAMs and hence the presence of nonlinear age effect.

2.11.2 Residual ComBat effects on real data.

Figure 2.9 presents the residual plots obtained after fitting a comprehensive model to the harmonized MRI-derived phenotype data, incorporating variables such as age, sex, and group. These residual plots provide insightful visual representations of the model's performance. Notably, the observed absence of pronounced shifts or variance scaling in all the models suggests that a linear correction approach proves to be more than adequate for addressing the particular use case (see Section 2.6).

Region	ΔAIC	Region	ΔAIC	Region	ΔAIC
rh_rostralmiddlefrontal_thickness	-737.827	rh_lateralorbitofrontal_thickness	-294.661	rh_fusiform_thickness	-128.352
Right-Hippocampus	-666.327	rh_superiorfrontal_thickness	-284.615	lh_caudalmiddlefrontal_thickness	-126.349
Left-VentralDC	-631.412	rh_lateraloccipital_thickness	-284.338	CC_Anterior	-124.647
Left-Hippocampus	-588.296	rh_posteriorcingulate_thickness	-277.702	lh_fusiform_thickness	-124.049
lh_precuneus_thickness	-586.472	rh_pericalcarine_thickness	-277.491	rh_isthmuscingulate_thickness	-111.719
Right-VentralDC	-577.552	CC_Mid_Posterior	-276.176	rh_parsopercularis_thickness	-106.362
lh_rostralmiddlefrontal_thickness	-570.793	lh_lateralorbitofrontal_thickness	-258.546	rh_parahippocampal_thickness	-105.893
lh_cuneus_thickness	-544.505	rh_parsorbitalis_thickness	-246.652	Right-Cerebellum-Cortex	-100.931
Left-Lateral-Ventricle	-530.619	rh_rostralanteriorcingulate_thickness	-243.986	Right-Pallidum	-94.757
rh_cuneus_thickness	-515.156	lh_parsorbitalis_thickness	-230.732	rh_middletemporal_thickness	-87.857
Right-Lateral-Ventricle	-514.104	Left-Caudate	-229.629	lh_parahippocampal_thickness	-83.225
rh_precuneus_thickness	-505.154	Right-Caudate	-226.743	lh_parsopercularis_thickness	-67.404
rh_lingual_thickness	-489.542	CC_Central	-217.088	lh_transversetemporal_thickness	-67.285
Brain-Stem	-451.604	lh_parstriangularis_thickness	-215.885	lh_superiortemporal_thickness	-63.792
lh_paracentral_thickness	-443.599	lh_lateraloccipital_thickness	-213.35	lh_middletemporal_thickness	-59.581
lh_lingual_thickness	-443.392	CC_Mid_Anterior	-209.57	rh_superiortemporal_thickness	-57.407
lh_inferiorparietal_thickness	-439.908	lh_superiorfrontal_thickness	-206.472	rh_transversetemporal_thickness	-56.076
lh_superiorparietal_thickness	-424.113	rh_caudalmiddlefrontal_thickness	-206.01	lh_insula_thickness	-54.416
rh_superiorparietal_thickness	-420.171	rh_entorhinal_thickness	-194.472	rh_insula_thickness	-53.683
rh_medialorbitofrontal_thickness	-417.312	lh_caudalanteriorcingulate_thickness	-191.647	lh_precentral_thickness	-51.688
Left-Amygdala	-409.543	lh_medialorbitofrontal_thickness	-189.768	lh_inferiortemporal_thickness	-48.254
Right-Cerebellum-White-Matter	-392.708	lh_entorhinal_thickness	-176.719	Left-Pallidum	-40.941
rh_inferiorparietal_thickness	-381.635	lh_isthmuscingulate_thickness	-168.909	rh_precentral_thickness	-37.279
rh_paracentral_thickness	-358.113	rh_supramarginal_thickness	-161.94	rh_inferiortemporal_thickness	-37.05
Right-Amygdala	-324.232	lh_postcentral_thickness	-157.198	Left-Putamen	-34.233
CC_Posterior	-323.508	rh_postcentral_thickness	-154.755	Right-Putamen	-24.207
rh_parstriangularis_thickness	-310.166	lh_supramarginal_thickness	-148.167		
lh_posteriorcingulate_thickness	-306.379	rh_caudalanteriorcingulate_thickness	-144.195		
lh_pericalcarine_thickness	-305.481	lh_rostralanteriorcingulate_thickness	-128.495		

Tab. 2.3: Difference of Akaike information criterion (AIC) between a nonlinear and a linear model using age as the exogenous variable across cortical thickness measures. $\Delta AIC < 0$ suggest that the nonlinear model is a best fit hence suggesting non linear effects of age in such negative differences.




Fig. 2.9: Extension of Figure 2.6 for all the methods evaluated in this work and top regions in Figure 2.5a.

2.12 Identifiability of ComBat parameters

A statistical model is said to be identifiable if there exists a unique set of parameter values that can explain the observed data. Mathematically, a statistical model is identifiable if the following conditions hold:

Let \mathcal{X} be the sample space and $f_\theta(x)$ be the probability density function or probability mass function of the statistical model, indexed by the parameter vector $\theta \in \Theta$.

1. θ is finite-dimensional and $\Theta \subseteq \mathbb{R}^d$ for some integer d .
2. For any $\theta_1, \theta_2 \in \Theta$, if $f_{\theta_1}(x) = f_{\theta_2}(x)$ for all $x \in \mathcal{X}$, then $\theta_1 = \theta_2$. In other words, the likelihood function, which is defined as $L(\theta; x) = f_\theta(x)$, is injective with respect to θ .

Condition 1 ensures that the parameter vector is finite-dimensional and belongs to some subset of Euclidean space. Condition 2 ensures that there exists a one-to-one correspondence between the parameter values and the probability distribution, which is necessary for unique estimation of the parameters. These one-to-one correspondence in the context of Fed-ComBat in Equation (2.3) ensures that the true expectation effect is captured by $\phi(\mathbf{x}_{ij}; \theta) + \alpha_g$ and not by γ_{ig} . Equivalently for Equation (2.4). In practice this is translated to:

- No allowing intercepts in the neural network architectures so only one is estimated by α_g .
- Using a design matrix \mathbf{z} shown in Equation (2.15) that allow the estimation of a single set of intercepts that can then be decomposed $\beta_g = \{\alpha_g + \gamma_{ig}\}_{i=1}^S$.

2.13 CRediT author statement

Santiago Silva: Conceptualization, Methodology, Software, Validation, Formal Analysis, Investigation, Writing - Original Draft, Visualization. **Neil Oxtoby:** Resources, Data Curation, Writing - Review & Editing. **Andre Altmann:** Supervision, Validation, Writing - Review & Editing. **Marco Lorenzi:** Supervision, Methodology, Validation, Writing - Review & Editing.

Federated black-box Bayesian optimization

3.1	Introduction	50
3.2	Overview of Black Box Optimization	51
3.2.1	Bayesian optimization	52
3.2.2	Acquisition functions	53
3.3	Methodology	55
3.3.1	Federated Bayesian Optimization	57
3.4	Results	59
3.4.1	Synthetic Data	59
3.4.2	Brain imaging data	61
3.4.3	Classification	62
3.4.4	Variational autoencoder	63
3.5	Conclusions	63

Abstract: One of the most critical steps in federated learning involves the iterative aggregation of models from multiple clients while optimizing the communication channel to reduce communication rounds and maintain generalization. In this work, we propose a black-box federated Bayesian optimization approach for model aggregation in federated learning, considering the learning of search space (bounds) for the parameters. Our method was evaluated on synthetic data and morphological measures from multi-centric databases for Alzheimer’s disease. We assessed the applicability of Bayesian optimization as part of the optimization scheme in federated learning on synthetic data and on multicentric brain data from four different studies. Results show that despite the promising potential of Bayesian optimization, standard approaches such as FedAVG are more suitable and scalable than Bayesian optimization, especially in complex problems where the number of parameters increases.

3.1 Introduction

Federated learning (FL) has emerged as a field that empowers institutions to participate in multicentric studies while preserving data governance. This collaborative learning paradigm allows multiple clients or data owners to share only the parameters of locally trained models with a central node, enabling larger sample access and encouraging consensus between clinical centers and researchers without exchanging or transferring data [Rieke et al., 2020].

Despite its potential, FL faces multiple challenges, including communication usage, hardware heterogeneity, and resource availability across institutions, which may prevent them from harnessing the power of their data. Additionally, ensuring the generalizability of models in heterogeneous setups, such as those found in collaborative studies, is crucial.

Existing federated optimization methods, like FedAVG [McMahan et al., 2017], FedProx [Li et al., 2018], and FedDANE [Li et al., 2019a], focus on aggregating locally partially trained models to increase generalizability. However, they often require multiple rounds of communication and depend on institutions having the computational power and resources to perform local gradient-based optimizations. This can be particularly challenging for institutions with limited resources or when facing hardware heterogeneity across different participants.

In contrast to the current federated optimization methods, centralized approaches like black-box optimization (BBO) and, more specifically, Bayesian optimization (BO) offer alternative solutions by building a response surface to optimize complex functions. BBO and BO optimize expensive-to-evaluate functions without relying on gradient descent, focusing on evaluating the target function instead. BO is particularly well-suited for tasks involving expensive-to-evaluate functions, as it offers a data-efficient and robust approach to global optimization.

Furthermore, BO provides convergence guarantees [Bull, 2011] and is especially suitable for problems with multiple minima due to its ability to construct a surrogate of the response surface. By leveraging the probabilistic representation of the objective function, BO not only accounts for the current best observation but also considers the uncertainty in the predictive distribution. Despite its use in centralized setups, BBO approaches have not been yet considered as an alternative for federated learning optimization.

In this work, we propose a federated optimization approach relying on black-box optimization that aims to centralize the optimization process in FL, reduce communication, and increase or be comparable in terms of generalizability by limiting participants of

collaborative studies to only evaluate the objective function without requiring local gradient optimization. We approximate a response surface with the results collected from the institutions, leveraging BO's probabilistic representation of the objective function, which not only accounts for the current best observation but also considers the uncertainty in the predictive distribution.

Our proposed method delegates the majority of the computational load to the central server, resulting in a centralized optimization process that is not affected by variations in available computing resources across different institutions. This solution addresses the challenges of communication usage, hardware heterogeneity, and resource availability while also ensuring the generalizability of models in heterogeneous setups. By doing so, we hope to pave the way for future research in federated learning, ultimately advancing the field and enabling more effective utilization of decentralized data in healthcare.

3.2 Overview of Black Box Optimization

Black box optimization (BBO) is a powerful approach for optimizing complex functions without any knowledge of their internal workings, relying solely on their evaluation at different input points [Bajaj et al., 2021]. This technique is particularly useful when the function is computationally expensive to evaluate, has an unknown analytical form, or has many parameters that need to be tuned. Black box optimization has found numerous applications in machine learning, where it is used for hyperparameter tuning [Bergstra et al., 2012] and neural architecture search [Real et al., 2019]. It is also widely used in engineering and scientific fields for optimizing expensive simulations and experiments [forrester2008engineering].

In black box optimization, the goal is to find the optimal set of input values that minimize or maximize the function of interest. This can be formulated as an optimization problem of the form:

$$\theta^* = \arg \min_{\theta \in \Omega} f(\theta), \quad (3.1)$$

where θ is a vector of input values, Ω is the set of feasible input values (search space), and $f(\theta)$ is the function to be optimized. The optimization problem in Equation 3.1 is often challenging to solve because the function $f(\theta)$ is typically complex, expensive to evaluate, and may have multiple local minima. Black box optimization methods aim to efficiently search the input space to find the global minimum or maximum of the function, without relying on any knowledge of its internal workings.

Several strategies exist for solving the optimization problem in Equation 3.1. For instance, Grid search exhaustively evaluates the function at all combinations of parameters in the search space, which is computationally expensive and may not scale well to high-dimensional search spaces. Random search samples parameters uniformly at random from the search space and evaluates the target function at these points, which is more efficient than grid search but may still require a large number of function evaluations [Bergstra et al., 2012].

Another popular approach to black box optimization is Bayesian optimization (BO), which uses probabilistic models to guide the search towards promising regions of the search space [shahriari2015taking]. This approach is especially useful when the objective function may not only be complex but also has a limited number of evaluations (i.e. a communication budget).

3.2.1 Bayesian optimization

Bayesian optimization is an optimization technique for expensive black-box functions, where the objective function is treated as a random function and modeled using Gaussian Processes (GPs) [MacKay et al., 2003]. It is particularly useful when evaluating the function is time-consuming or costly, and the number of evaluations is limited. Now the target function is consider as:

$$f(\boldsymbol{\theta}) \sim \mathcal{GP}(\mu(\boldsymbol{\theta}), k(\boldsymbol{\theta}, \boldsymbol{\theta}')), \quad (3.2)$$

where $f(\boldsymbol{\theta})$ represents the surrogate model of the objective function, approximated by a GP with mean function $\mu(\boldsymbol{\theta})$ and covariance function $k(\boldsymbol{\theta}, \boldsymbol{\theta}')$. With its goal being to optimize the problem in Equation (3.1), the use of a GP as a model of the target function requires a guide for searching the optimal candidate parameters at each iteration. This guide is a function known as an “acquisition function”.

Acquisition functions guide the search for the optimal input by balancing exploration and exploitation. They quantify the utility of evaluating the objective function at a particular input point, given the current model. Probability of Improvement (PI), Expected Improvement (EI), and Upper Confidence Bound (UCB) are some of the popular acquisition functions used in Bayesian optimization [Brochu et al., 2010]. These are described more in detail in Section 3.2.2.

The optimization process using BO takes place iteratively as the response surface (target function) is approximated by the surrogate model (GP) and new samples are proposed as follows:

1. Optimize the acquisition function over the GP to find the next sampling point in the search space: $\theta^t = \arg \max_{\theta \in \Omega} a(f(\theta))$. Where a denotes the acquisition function.
2. Compute a noisy evaluation of the black-box function using the suggested set of input values from step 1, θ^t : $y^t = f(\theta^t) + \epsilon_t$.
3. Update the GP with the new known sample from the search space and its measured response (e.g. loss value): $\Omega = \Omega \cup \{\theta^t\}$ and $\mathcal{Y} = \mathcal{Y} \cup \{y^t\}$.

This is repeated until a stopping criterion is met, such as a maximum number of iterations or a convergence threshold. The optimal θ^* will be finally defined as the set of input values that minimize the GP model.

3.2.2 Acquisition functions

The acquisition function serves as a guide to balance *exploration* and *exploitation* in BO. Its main objective is to guide the search for the optimum by directing it towards regions where the objective function is expected to be minimized/optimized. This can be accomplished either by selecting points with low predictions based on the probabilistic surrogate model or by targeting regions with high uncertainty in the model's predictions (or both).

Exploration refers to the act of seeking out and gathering information about new or unexplored regions of the search space. It involves trying out different options and sampling diverse areas to gain a broader understanding of the problem. The goal of exploration is to acquire new knowledge, uncover potential opportunities, or identify previously unknown solutions.

Exploitation, on the other hand, involves maximizing the current knowledge or exploiting known resources to achieve immediate gains. It focuses on utilizing the existing information, or exploiting well-established solutions to optimize the outcome of the objective function. The aim of exploitation is to capitalize on the available knowledge to obtain the best possible results based on the current understanding of the problem.

Among some of the most established acquisition functions there is:

Probability of Improvement (PI)

Probability of Improvement (PI) is an acquisition function that selects the next point for evaluation based on the probability that it will improve over the current best observation [Jones et al., 1998a]. Mathematically, PI can be defined as:

$$\text{PI}(\boldsymbol{\theta}) = P(f(\boldsymbol{\theta}) > f(\boldsymbol{\theta}^+) + \xi), \quad (3.3)$$

where $f(\boldsymbol{\theta})$ is the objective function, $\boldsymbol{\theta}^+$ is the current best observation, and ξ is an exploration parameter that encourages sampling in uncertain regions. The probability can be computed using the Gaussian process model's predictive distribution:

$$\text{PI}(\boldsymbol{\theta}) = \Phi\left(\frac{\mu(\boldsymbol{\theta}) - f(\boldsymbol{\theta}^+) - \xi}{\sigma(\boldsymbol{\theta})}\right), \quad (3.4)$$

where Φ is the cumulative distribution function of the standard normal distribution, $\mu(x)$ is the predictive mean, and $\sigma(\boldsymbol{\theta})$ is the predictive standard deviation at the input point x [Mockus, 1998].

Expected Improvement (EI)

Expected Improvement (EI) is another acquisition function that balances exploration and exploitation by considering the expected improvement over the current best observation [Mockus, 1998]. EI can be defined as:

$$\text{EI}(x) = E\left[\max(f(\boldsymbol{\theta}) - f(\boldsymbol{\theta}^+), 0)\right], \quad (3.5)$$

where $f(\boldsymbol{\theta})$ is the objective function, and $\boldsymbol{\theta}^+$ is the current best observation. Using the Gaussian process model's predictive distribution, the expected improvement can be computed as:

$$\text{EI}(\boldsymbol{\theta}) = (\mu(\boldsymbol{\theta}) - f(\boldsymbol{\theta}^+) - \xi)\Phi(Z) + \sigma(\boldsymbol{\theta})\phi(Z), \quad (3.6)$$

where $Z = \frac{\mu(\theta) - f(\theta^+) - \xi}{\sigma(\theta)}$, Φ is the cumulative distribution function, and ϕ is the probability density function of the standard normal distribution. The parameter ξ encourages exploration, with larger values promoting greater exploration of the search space [Jones et al., 1998b].

Upper Confidence Bound (UCB)

Upper Confidence Bound (UCB) is an acquisition function that selects the next point for evaluation based on the upper confidence bound of the predictive distribution [Auer, 2002]. UCB can be defined as:

$$\text{UCB}(\theta) = \mu(\theta) + \kappa\sigma(\theta), \quad (3.7)$$

where $\mu(\theta)$ is the predictive mean, $\sigma(\theta)$ is the predictive standard deviation at the input point θ , and κ is a parameter that controls the exploration-exploitation trade-off. A larger value of κ encourages more exploration.

3.3 Methodology

In a federated scenario, we suppose there are K institutions participating in a federated learning experiment, each with a local dataset \mathcal{D}_k of size N_k . The goal is to learn a global model θ that minimizes the expected loss over all clients, subject to privacy and communication constraints. This can be formulated as the following optimization problem:

$$\min_{\theta \in \Omega} f(\theta) + \lambda \mathcal{R}(\theta), \quad (3.8)$$

where,

$$f(\theta) = \sum_{k=1}^K \frac{N_k}{N} F_k(\theta), \quad (3.9)$$

Ω is the set of feasible parameters for the model, $N = \sum_{k=1}^K N_k$ is the total number of data points, $F_k(w)$ is the local loss function for client k that measures the difference between the predicted output of the model θ and the true output for a given input x , and $\mathcal{R}(\theta)$ is a regularization term that encourages the model to be simple or have

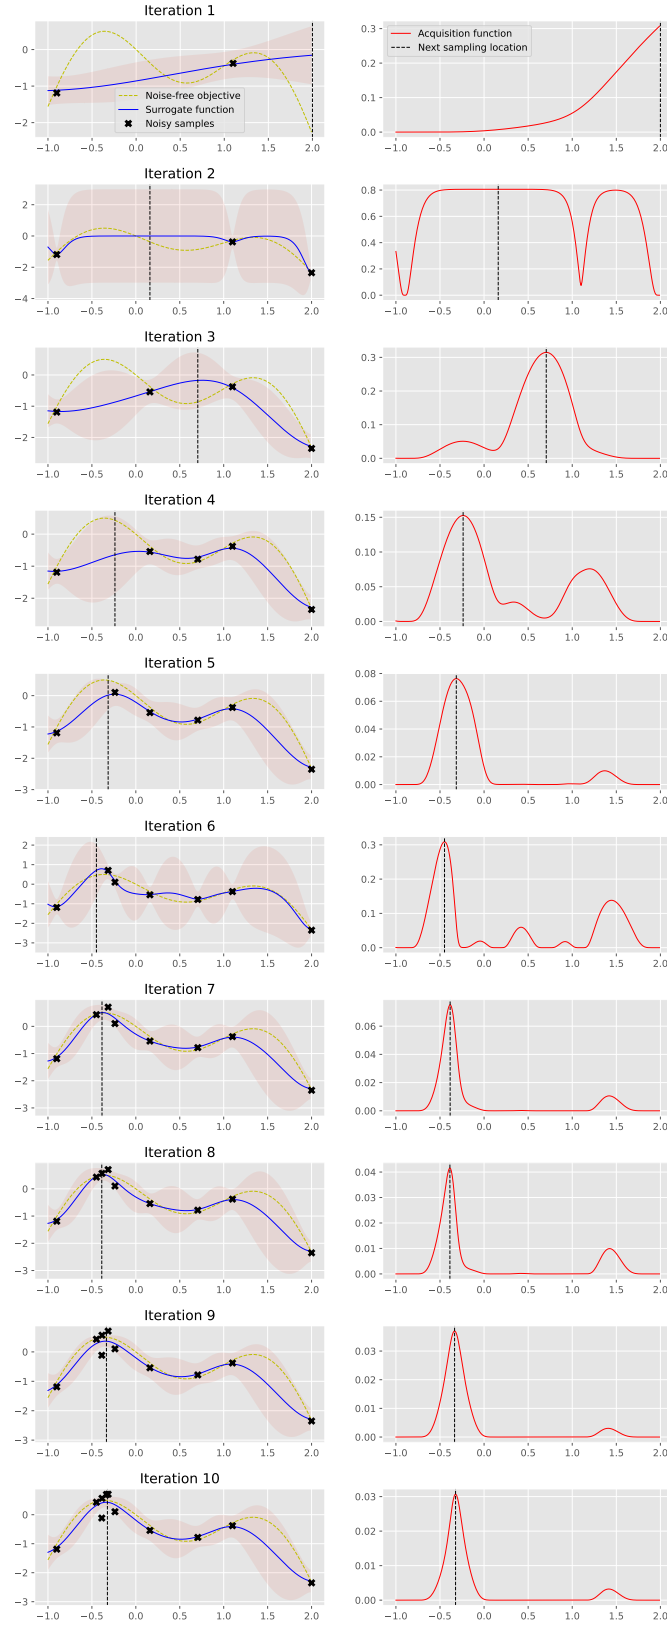


Fig. 3.1: Bayesian optimization procedure with 10 iterations. Each row of two plots are result of an iteration. The left plot depicts the noise-free objective function alongside the surrogate function, represented by the Gaussian process posterior predictive mean. Additionally, the 95% confidence interval of the mean and the noisy samples obtained from the objective function are illustrated. The right plot showcases the acquisition function, specifically the expected improvement. Notably, a vertical dashed line is included in both plots to indicate the proposed sampling point for the subsequent iteration, which corresponds to the maximum value of the acquisition function.

certain desirable properties, and λ is a hyperparameter that controls the strength of the regularization.

The local loss function $F_k(w)$ for client k is defined as:

$$F_k(\boldsymbol{\theta}) = \frac{1}{N_k} \sum_{(x,y) \in D_k} \ell(\boldsymbol{\theta}; x, y) \quad (3.10)$$

where $\ell(w; x, y)$ is the loss function that measures the difference between the predicted output of the model w for a given input x and the true output y .

We propose to solve this optimization problem by iteratively updating the global model $\boldsymbol{\theta}$ using local evaluations of the target function from each client without performing local updates as originally proposed by [McMahan et al., 2017]. Instead of aggregating the locally updated models at each step, a response surface is better approximated at each iteration by a GP as explained in the next section. The process is repeated until convergence or a stopping criterion is met.

In our approach, the regularization term $\mathcal{R}(\boldsymbol{\theta})$ serves a similar purpose to the exploration-exploitation trade-off terms mentioned earlier. It is a function of the acquisition function and helps balance the exploration of new regions of the input space with the exploitation of regions with high expected utility avoiding thus falling in local minima.

3.3.1 Federated Bayesian Optimization

In our approach, by considering each client's objective functions $F_k(\boldsymbol{\theta})$ as realizations of a Gaussian process as shown in Equation (3.11), then minimizing the GP with respect to the parameters of the model is equivalent to finding the best operating parameters $\boldsymbol{\theta}$ of the model in question. In order to accelerate convergence we conveniently consider that $\boldsymbol{\theta}$ lives in a subspace or parameters denominated the *search space* Ω .

$$F_k(\boldsymbol{\theta}) \sim \mathcal{GP}(\mu(\boldsymbol{\theta}), k(\boldsymbol{\theta}, \boldsymbol{\theta}')) \implies \min_{\boldsymbol{\theta} \in \Omega} f(\boldsymbol{\theta}) \propto \min_{\boldsymbol{\theta} \in \Omega} \mathcal{GP}(\mu(\boldsymbol{\theta}), k(\boldsymbol{\theta}, \boldsymbol{\theta}')) \quad (3.11)$$

Search space learning:

Accelerating the convergence of optimization can be achieved by defining appropriate bounds for the search space. Learning the search space based on previous evaluations of optimized black-box functions has been proposed by [Perrone et al., 2019]. In the context of federated learning, the search space can be defined using the results of locally optimized models from each participating institution.

Valerio et al. formulate the problem as minimizing the volume of the search space \mathcal{Q} , determined by the bounds of the subspace Ω . The goal is to ensure that all the solutions resulting from maximizing the acquisition function θ^+ are contained within Ω , as shown in Equation (3.12). By learning the search space geometry from historical data, this approach can endow many existing Bayesian optimization methods with transfer learning properties, leading to a more efficient optimization process.

$$\min_{\Omega \in \mathbb{R}^d} \mathcal{Q}(\Omega) \text{ such that for } t \geq 1, \theta^+ \in \Omega \quad (3.12)$$

As a first instantiation, considering the bounding space a hyperrectangle, whose search space can be described by its lower and upper bounds $\Omega = (\mathbf{l}, \mathbf{u})$, the problem can then be rewritten as:

$$\min_{\mathbf{l} \in \mathbb{R}^d, \mathbf{u} \in \mathbb{R}^d} \frac{1}{2} \|\mathbf{u} - \mathbf{l}\|_2^2 \quad \text{such that,} \quad \mathbf{l} \leq \theta^+ \leq \mathbf{u} \quad (3.13)$$

$$\text{Hence, } \mathbf{l}^* = \min\{\theta^+\}_{t=1}^T \quad \text{and} \quad \mathbf{u}^* = \max\{\theta^+\}_{t=1}^T \quad (3.14)$$

Optimization of the target function:

We rely on BO and the probabilistic nature of GPs to estimate not only the mean but also the uncertainty of a function. We choose Expected Improvement (EI) as the acquisition function, as it is considered less prone to falling into local minima. The fundamental idea behind EI is to maximize the distance between the current point and the candidate point, while balancing the trade-off between exploitation and exploration. EI focuses on not only identifying candidates at the optima of $\mu(\theta)$ but also penalizing the uncertainty $\sigma^2(\theta) = k(\theta, \theta) - k(\theta, \{\theta^+\}_{t=1}^T) [K(\{\theta^+\}_{t=1}^T, \{\theta^+\}_{t=1}^T) + \lambda^2 I]^{-1} k(\{\theta^+\}_{t=1}^T, \theta)$ [Pardalos et al., 2021]. This balancing act enables a trade-off between exploitation (finding the minimum) and exploration (avoiding local optima). EI, in particular, evaluates the expected positive distance between the current point and the solution θ^+ , offering an effective approach to guide the optimization process.

$$\text{EI}(\theta) = \mathbb{E} \max(f(\theta) - f(\theta^+), 0) \quad (3.15)$$

$$\text{EI}(\theta) = \begin{cases} (\mu(\theta) - f(\theta^+) - \xi)\Phi(Z) + \sigma(\theta)\phi(Z) & \text{if } \sigma(\theta) > 0 \\ 0 & \text{if } \sigma(\theta) = 0 \end{cases} \quad (3.16)$$

$$Z = \begin{cases} \frac{\mu(\theta) - f(\theta^+) - \xi}{\sigma(\theta)} & \text{if } \sigma(\theta) > 0 \\ 0 & \text{if } \sigma(\theta) = 0 \end{cases} \quad (3.17)$$

Where, $\mu(\theta)$ and $\sigma(\theta)$ correspond to the posteriori prediction from the \mathcal{GP} . Φ and ϕ are the CDF and PDF respectively. The parameter ξ is a penalisation term accounting for exploration in the parameter space.

Each solution θ^+ in this context is somehow an equivalent to the aggregated model deferring from it that it is meant to explore the search space to gather as much knowledge as possible from the response function f in Equation (3.9) instead of advancing in the direction with less gradient.

The aggregation scheme relying on BO and the learning bounds strategy of [Perrone et al., 2019] is explained in ?? 2.

Algorithm 2: Federated Bayesian optimisation

Require: Define a model \mathcal{M}_θ , parametrized by θ , and an objective function

$f(\theta) \sim \mathcal{GP}(\mu(\theta), k(\theta, \theta'))$. N_R rounds and K clients.

Ensure: Initialize set of parameters: $\theta^+ \sim \text{Uniform}(\mathbf{l}, \mathbf{u})$.

```

while  $n_i \leq N_R$  do
  for  $k \in \{1, 2, \dots, K\}$  do
    Evaluate  $F_k(\theta^+)$ 
  end for
  Compute global loss  $y^t \leftarrow \sum_{k=1}^K \frac{N_k}{N} F_k(\theta^+)$ 
  Update explored space  $\Omega \leftarrow \Omega \cup \theta^+, \mathbf{y} \leftarrow \mathbf{y} \cup y^t$ 
  Fit response surface  $f(\theta) \leftarrow \mathcal{GP}(\cdot, \cdot) | (\mathbf{y}, \Omega)$ 
  Obtain  $\theta^+ = \max_w \text{EI}(\theta)$ 
end while

```

3.4 Results

3.4.1 Synthetic Data

Synthetic data was generated for two tasks: regression, and classification. These two tasks were limited to linear regression and logistic classification on synthetic data as these are simple problems and convex in the case of linear regression.

Scalability

We tested for 5 different levels of complexity (numbers of parameters): 1, 3, 10, 30, and 100. The data was slightly biased per site with a random shift, in order to simulate a more realistic scenario. This test of dimensionality allows to compare the scalability of FedAVG vs our proposed Bayesian optimization method.

As shown in Figure 3.2, the scalability of FedAVG and our approach was compared for linear regression with varying numbers of parameters. The results indicate that as the

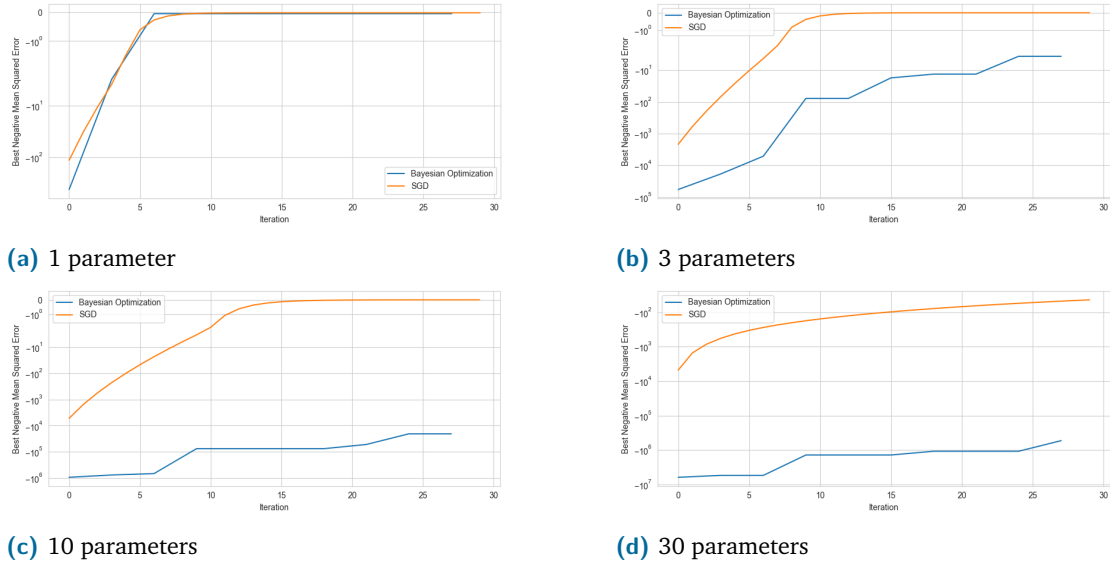


Fig. 3.2: Negative mean squared error over iterations for different number of parameters on synthetic data. Model optimized: Linear regression.

number of parameters increases, FedAVG scales better than our approach. Figure 3.3 illustrates the scalability of both approaches with respect to the groundtruth set of parameters and the execution time as the number of parameters increases logarithmically.




Fig. 3.3: Euclidean and cosine distance between optimized and true parameters using FedAVG and our proposal across several number of parameters. Lower is better.

Additionally to the adaptive bounds we tested a combination of locally optimized using SGD for one epoch to better guide the search of the search of the optima. An example of it is shown in Figure 3.4. The result drastically increases as a more guided exploration of the parameter space is done by gradient descent and thus leading to a better convergence. However, the optimization does not improve after given that it keeps exploring instead of exploiting. Meaning most of the optimization is done by gradient descent itself.

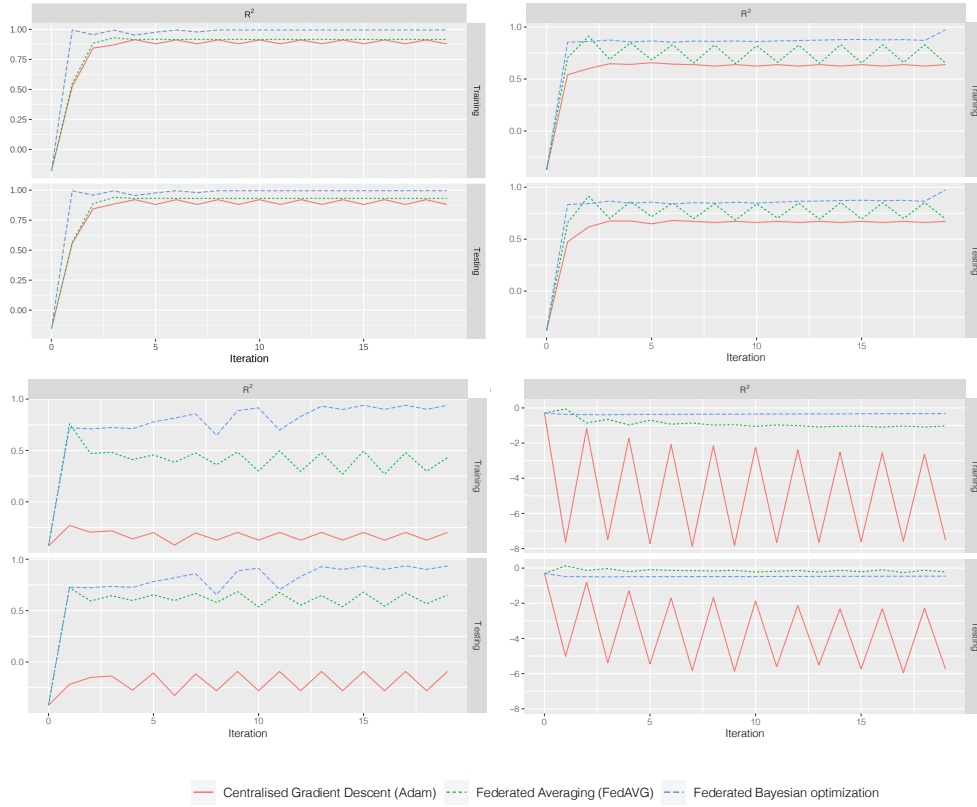


Fig. 3.4: Evolution of R^2 for regression in a convex problem on synthetic data across different scenarios varying the number of parameters to optimize. FedAVG settings: 10 epochs locally. Each iteration corresponds to a round of communication/aggregation. Parameters being optimised: 30 (top left), 100 (top right), 300 (bottom left) and 1000 (bottom right).

An example on more complex model optimization such as logistic classifiers and variational autoencoders is shown on left panels in Figure 3.5 and Figure 3.6 respectively.

3.4.2 Brain imaging data

A total of 87 measures consisting of subcortical volumes and cortical thicknesses extracted from T1 weighted magnetic resonance images (T1w-MRI) using FreeSurfer 6.0. Different datasets were used in order to have a non-iid multi-center realistic scenario: 801 participants at baseline from the Alzheimer’s Disease Neuroimaging Initiative (ADNI), 401 participants from the OASIS-I cross-sectional dataset, 1117 from OASIS-III, and 67 from the Minimal Interval Resonance Imaging in Alzheimer’s Disease (MIRIAD). Groups were classified by controls (NC) and cognitively impaired (CI) for classification. Extensive description of population’s demographics is shown in table 5.2.

		Controls (NC)	Cognitively Impaired (CI)
ADNI (N=801)	N. Participants	175	621
	Age at study entry (years \pm std)	75.79 ± 4.99	74.86 ± 7.23
	Men	52%	58%
OASIS-I (N=401)	N. Participants	314	87
	Age at study entry (years \pm std)	54.22 ± 26.20	78.86 ± 11.12
	Men	37%	41%
OASIS-III (N=1098)	N. Participants	709	306
	Age at study entry (years \pm std)	67.8 ± 9.86	74.28 ± 7.87
	Men	41%	53%
MIRIAD (N=69)	N. Participants	23	46
	Age at study entry (years \pm std)	69.7 ± 7.2	69.4 ± 7.1
	Men	52%	41%

Tab. 3.1: Demographics from all the datasets used in this work.

A classification task between case-control was performed using a logistic classifier guiding the search with one epoch of SGD locally. Evolution of the classification metrics is illustrated on the right panel of Figure 3.5.

3.4.3 Classification

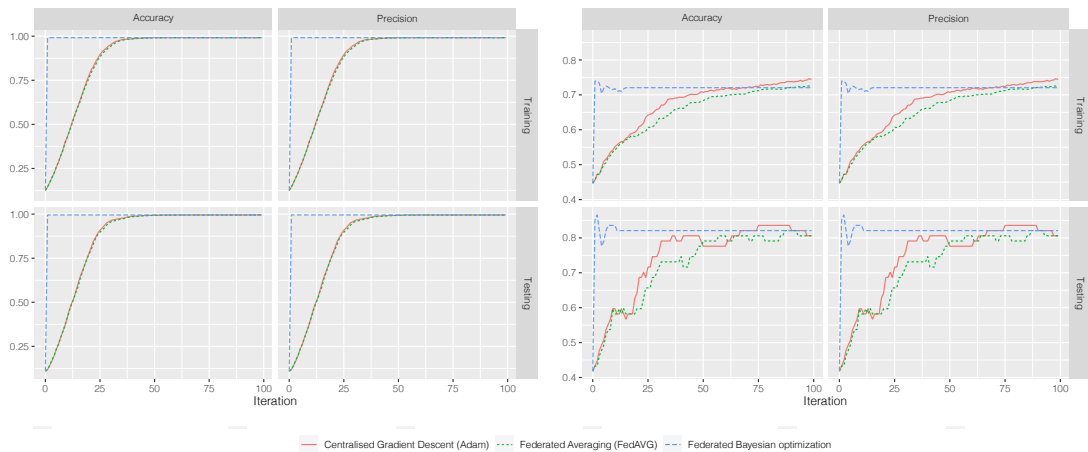


Fig. 3.5: Evolution of classification metrics for synthetic data (400 parameters and 4 classes) evolution on training and test sets sampled and stratified by class (left). Results for brain data using MIRIAD only as testing set (right). FedAVG settings: 10 epochs locally.

Equivalently, results on the reconstruction error for a variational autoencoder are illustrated on the right panel of Figure 3.6.

3.4.4 Variational autoencoder

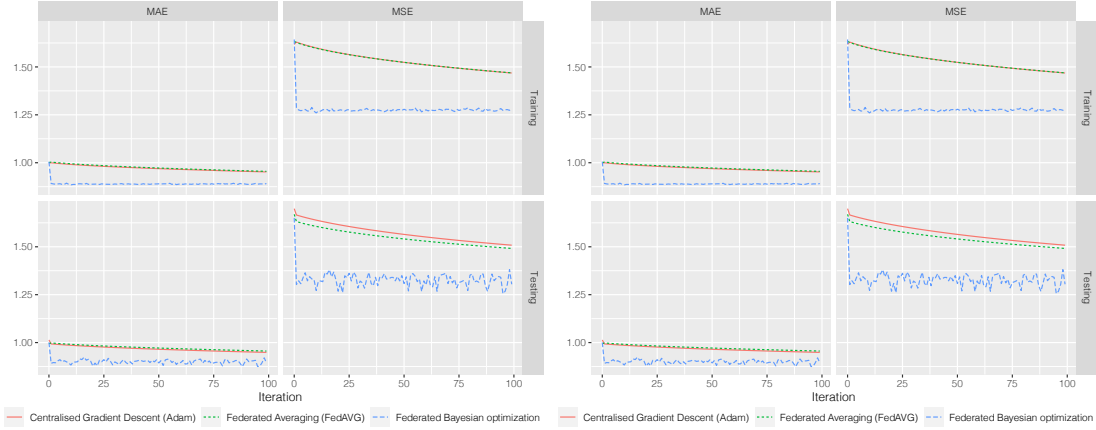


Fig. 3.6: Evolution of reconstruction losses (MAE and MSE) for Variational autoencoder on synthetic data (left) coming from 2 classes and a latent space equal to the number of classes. Same results applied on brain morphological data (right). FedAVG settings: 10 epochs locally.

3.5 Conclusions

We investigated in this chapter the use of Bayesian optimization as an aggregation scheme for federated learning. While it shows promise in centralizing the optimization process and potentially improving the generalizability of models, scalability remains a challenge. The number of parameters in a model impacts the optimization process, even when adaptive bounds are applied. One potential solution to address this issue is to guide the initial search using gradient descent. However, further exploration is needed, as our findings indicate that the model's performance does not improve substantially after the initial search. In conclusion, Bayesian optimization offers a promising direction for federated learning, but additional research is required to address the scalability challenges and fully harness its potential in this context.

Federated Data Harmonization in Biomedical Research using Mixed Effects Models: A Focus on Conditional Variational Autoencoders

4.1	Introduction	65
4.2	Fundamentals of Data Harmonization	66
4.2.1	Significance of Data Harmonization in Biomedical Research	67
4.2.2	Challenges and Issues in Data Harmonization	67
4.3	Methods	68
4.3.1	Federated Harmonization with Conditional Variational Autoencoders (CVAEs)	68
4.3.2	Harmonization	69
4.4	Results	70
4.5	Challenges and Future Directions	71
4.6	Conclusions	75

4.1 Introduction

Biomedical research has rapidly evolved in recent years, thanks to advances in high-throughput technologies, which have generated a wealth of heterogeneous and multidimensional data. Data harmonization, the process of integrating and reconciling data from diverse sources and formats, has become a crucial task in the field [fortin2018data].

However, in multicentric neuroimaging studies, bias is introduced due to different acquisition protocols, magnetic resonance imaging (MRI) scanner brands, and other factors [Orlhac et al., 2022]. This bias can significantly affect the generalizability of the derived models and limit the potential of collaborative research. Several harmonization techniques have been developed to address this bias, including surrogate variable analysis [Leek et al., 2012], ComBat [Johnson et al., 2007] and random effect models [Kim et al., 2022]. However such methods remain univariate and are not able to properly capture

more complex relationships like nonlinearities. Also, these methods currently rely on data being centralized which as discussed, is conflicting with the current data protection regulations.

Federated learning (FL) is an emerging paradigm in the field of machine learning that enables multicentric studies in brain imaging data to remain compliant with data protection regulations such as the General Data Protection Regulation (GDPR) and California Consumer Privacy Act (CCPA). By distributing the learning process across multiple institutions while keeping the data locally, FL ensures better data governance and security [Li et al., 2020a].

Existing methods such as NeuroHarmony [Garcia-Dias et al., 2020] and domain adaptation methods on raw imaging data [Zuo et al., 2021] also have limitations. NeuroHarmony, which is based on NeuroComBat, considers only linear shifts and scaling, while Zuo's method depends on convolutional layers that may introduce artifacts if not properly trained. Additionally, these methods often require retraining when a new institution joins the study.

The aim of this chapter is to explore conditional generative models such as variational autoencoders (VAEs) for data harmonization in federated learning and their potential to address the limitations of existing techniques by capturing complex relationships and interactions while accounting for mixed effects introduced by different sites or institutions.

The chapter is organized as follows: Section 2 provides a background on the importance of data harmonization in biomedical research, while Section 3 introduces the concept of CVAEs and their potential for addressing data harmonization challenges. Section 4 delves into the applications of CVAEs in data harmonization, and Section 5 presents a detailed methodology for implementing CVAEs in this context. Section 6 discusses the results and performance evaluation of CVAEs in data harmonization, followed by an exploration of the challenges and future directions in Section 7. The chapter concludes with a summary of the key findings in Section 8.

4.2 Fundamentals of Data Harmonization

Data harmonization is a process that aims to create consistency and compatibility between datasets originating from different sources, formats, and measurement units, thereby enabling researchers to analyze and draw meaningful conclusions from the combined data. The need for data harmonization arises from the growing volume and complexity of biomedical data, which often exhibit heterogeneity due to variations in data collection

methodologies, experimental conditions, and measurement units. In this section, we discuss the significance of data harmonization in biomedical research, as well as the common challenges and issues associated with harmonizing heterogeneous datasets.

4.2.1 Significance of Data Harmonization in Biomedical Research

Data harmonization plays a critical role in biomedical research, as it allows for the integration and comparison of data from different sources, ultimately enhancing the reproducibility and generalizability of research findings [Fortin et al., 2014]. Furthermore, data harmonization can facilitate the discovery of novel associations and patterns in data, leading to new insights and hypotheses in biomedical research. By enabling researchers to leverage the full potential of the available data, data harmonization contributes to the acceleration of scientific discovery and the improvement of patient care.

4.2.2 Challenges and Issues in Data Harmonization

There are several challenges associated with data harmonization in biomedical research, including:

1. **Heterogeneity:** Datasets collected from different sources often exhibit variations in data formats, measurement units, and experimental conditions, making it difficult to directly compare and integrate the data.
2. **Missing data:** Incomplete or missing data points can hinder the harmonization process, as they may introduce bias and compromise the validity of the analysis.
3. **High dimensionality:** Biomedical datasets often consist of a large number of variables, which can pose computational challenges and increase the risk of overfitting during the harmonization process.
4. **Measurement errors:** Variations in data collection protocols, instrumentation, and human error can introduce measurement errors that need to be addressed during the harmonization process.
5. **Privacy and ethical considerations:** Harmonizing data from multiple sources raises privacy and ethical concerns, as it involves the integration and sharing of sensitive patient information.

4.3 Methods

4.3.1 Federated Harmonization with Conditional Variational Autoencoders (CVAEs)

To address the challenges of data harmonization in a federated setting, we propose a federated harmonization framework based on Conditional Variational Autoencoders (CVAEs). CVAEs have shown promise in capturing complex relationships and generating harmonized data samples conditioned on specific attributes or features.

Let \mathbf{y} denote the site indicator variable, z represent the latent space of the autoencoder, \mathbf{x} a design matrix containing the measured phenotypes and the covariates to be preserved, and θ and ϕ parametrize the autoencoder. We model θ and ϕ as random effects, centered at zero and with a variance σ^2 . We begin by defining the notation and deriving the evidence lower bound (ELBO).

To derive the ELBO, we assume that the true posterior distribution of the latent space, $p(z|\mathbf{x}, \mathbf{y})$, is intractable. We then introduce a variational approximation, $q_\phi(z|\mathbf{x}, \mathbf{y})$, which is parametrized by ϕ :

$$\log p(\mathbf{x}|\mathbf{y}) = \log \int p(\mathbf{x}|z, \mathbf{y})p(z|\mathbf{x}, \mathbf{y})dz \geq \mathbb{E}q_\phi(z|\mathbf{x}, \mathbf{y}) \left[\log \frac{p(\mathbf{x}|z, \mathbf{y})p(z|\mathbf{y})}{q_\phi(z|\mathbf{x}, \mathbf{y})} \right] \quad (4.1)$$

$$= \mathcal{L}(\theta, \phi; \mathbf{x}, \mathbf{y}) \quad (4.2)$$

The ELBO, $\mathcal{L}(\theta, \phi; \mathbf{x}, \mathbf{y})$, can be further decomposed into two terms: the reconstruction error and the Kullback-Leibler (KL) divergence between the variational approximation and the prior on the latent space:

$$\mathcal{L}(\theta, \phi; \mathbf{x}, \mathbf{y}) = \mathbb{E}q_\phi(z|\mathbf{x}, \mathbf{y})[\log p(\mathbf{x}|z, \mathbf{y})] - D_{KL}(q_\phi(z|\mathbf{x}, \mathbf{y})||p(z|\mathbf{y})) \quad (4.3)$$

$$= -\mathcal{F}(\theta, \phi; \mathbf{x}, \mathbf{y}) + \mathcal{D}(\phi; \mathbf{x}, \mathbf{y}) \quad (4.4)$$

We rely on stochastic variational inference (SVI) for the optimization scheme, which approximates the true posterior by minimizing the negative ELBO. In the federated setting, we employ federated averaging to optimize the global parameters. This involves aggregating the local updates of the variational parameters θ and ϕ across sites, weighted

by the number of samples at each site. This approach allows for the harmonization of multicentric data while addressing the privacy and regulatory concerns associated with sharing raw data between institutions.

The derivation of the lower bound generation using variational inference:

$$\log p(\mathbf{x}|\mathbf{y}) = \log \int p(\mathbf{x}|z, \mathbf{y})p(z|\mathbf{x}, \mathbf{y})dz \quad (4.5)$$

$$= \log \int p(\mathbf{x}|z, \mathbf{y}) \frac{p(z|\mathbf{x}, \mathbf{y})q_\phi(z|\mathbf{x}, \mathbf{y})}{q_\phi(z|\mathbf{x}, \mathbf{y})} dz \quad (4.6)$$

$$\geq \int q_\phi(z|\mathbf{x}, \mathbf{y}) \log \frac{p(\mathbf{x}|z, \mathbf{y})p(z|\mathbf{y})}{q_\phi(z|\mathbf{x}, \mathbf{y})} dz \quad (4.7)$$

$$= \mathbb{E}_{q_\phi(z|\mathbf{x}, \mathbf{y})} \left[\log \frac{p(\mathbf{x}|z, \mathbf{y})p(z|\mathbf{y})}{q_\phi(z|\mathbf{x}, \mathbf{y})} \right] \quad (4.8)$$

$$= \mathcal{L}(\theta, \phi; \mathbf{x}, \mathbf{y}) \quad (4.9)$$

Extending this formulation to a random effects model in which θ and ϕ are realizations of a distribution $p(\theta, \phi) = \mathcal{N}(0, \sigma^2)$, we can include their prior distributions in the optimization process. The updated ELBO with the prior distribution on the parameters θ and ϕ can be written as:

$$\begin{aligned} \mathcal{L}(\theta, \phi; \mathbf{x}, \mathbf{y}) &= \mathbb{E}_{q_\phi(z|\mathbf{x}, \mathbf{y})} [\log p(\mathbf{x}|z, \mathbf{y})] - \\ &\quad D_{KL}(q_\phi(z|\mathbf{x}, \mathbf{y}) || p(z|\mathbf{y})) - D_{KL}(q(\theta, \phi) || p(\theta, \phi)) \quad (4.10) \\ &= -\mathcal{F}(\theta, \phi; \mathbf{x}, \mathbf{y}) + \mathcal{D}(\phi; \mathbf{x}, \mathbf{y}) - \mathcal{R}(\theta, \phi) \end{aligned}$$

where $\mathcal{R}(\theta, \phi)$ denotes the regularization term due to the prior distribution on θ and ϕ . The learning of θ and ϕ can be achieved by minimizing the negative ELBO with respect to the parameters. In the federated setting, we employ federated averaging to optimize the global parameters, which involves aggregating the local updates of the variational parameters θ and ϕ across sites, weighted by the number of samples at each site.

4.3.2 Harmonization

Once the CVAE model has been trained using a standard federated learning approach such as FedAVG, we proceed to the harmonization step, which aims to align and reconcile the distributions of the data across different institutions. This step is essential for ensuring data consistency and comparability, enabling meaningful analysis and modeling within the federated learning framework.

We aim to address any remaining batch effects captured by the parameters of the CVAE. To achieve this, we leverage on the site indicator variable y , which represents the source institution or batch information for each data point. By sampling y from a multinomial distribution, in expectation, we can effectively average out the batch effects and reduce their influence on the reconstruction process.

Sampling y in expectation ensures that we obtain a diverse set of site indicator values for each observation x . By considering multiple samples of y for each data point, we capture the variability and distribution of the site indicator information. This averaging process helps to minimize the impact of any specific batch effect on the reconstruction and leads to a less biased representation of the underlying features.

By incorporating the expected values of y during the reconstruction step, we mitigate the influence of batch effects, allowing for a more accurate and unbiased reconstruction of the original data x . This approach improves the harmonization process by reducing the residual variability associated with different institutions or batches, ultimately leading to a more consistent and reliable representation of the underlying features across the dataset.

4.4 Results

We designed a synthetic dataset to evaluate the performance of the proposed harmonization method. The dataset incorporates covariates such as Sex, Age, and Scanner effects, which are commonly encountered in neuroimaging studies. These covariates capture important factors that can introduce variations in brain measures across different data sources or institutions.

By using principal component analysis (PCA) and analysis of variance (ANOVA), we gain insights into the effectiveness of the data harmonization method. Figure ?? showcases the results obtained through PCA. The top plot depicts the PCA of the non-harmonized features, where distinct clusters corresponding to different sites or institutions are evident. These clusters highlight the presence of batch effects, indicating potential variations and biases in the data.

After applying the data harmonization method, as illustrated in the middle plot, the feature distributions become more aligned, reducing the variations between sites and leading to a more integrated representation of the data. The clusters become less distinct, suggesting that the harmonization process successfully mitigates the impact of batch effects and promotes a more unified view of the dataset.

To further assess the impact of harmonization, ANOVA is conducted on the harmonized data. The results reveal a significant reduction in the variability attributed to the site indicator variable, indicating a successful harmonization outcome. This reduction in site-specific variability supports the notion that the data harmonization approach effectively mitigates batch effects and enables a more reliable and unbiased analysis of the dataset.


Figure 4.2 provides an overview of the feature distributions across multiple sites or institutions. The top plot reveals distinct clusters corresponding to different sites, indicating the presence of batch effects and potential biases in the data. In the middle plot, the feature distribution after harmonization is shown, demonstrating a reduction in variations between sites. This alignment of feature distributions aims to achieve a more consistent representation of the underlying features, effectively mitigating the impact of batch effects. The bottom plot represents the desired outcome: an unbiased distribution where batch effects have been entirely eliminated. This homogeneous feature distribution enables the collective analysis of data from different institutions, free from the confounding effects of batch variations.

We evaluated the effectiveness of the harmonization method using Bland-Altman plots. Figure 4.3 presents the feature differences between the unbiased dataset and the unharmonized dataset (Figure 4.3a), as well as between the unbiased dataset and the harmonized dataset using the CVAE method (Figure 4.3b), stratified by batch effects. In these plots, the x-axis represents the mean of the feature values, while the y-axis shows the difference between the corresponding feature values. The colored dots indicate different batches, enabling the assessment of batch effects on feature differences. Comparing the two plots, we observe a substantial reduction in feature differences in the harmonized dataset compared to the unharmonized dataset. This finding highlights the efficacy of the CVAE-based harmonization approach in mitigating the confounding effects of batch variations, thereby achieving a more consistent representation of the underlying features across diverse institutions or sites.


4.5 Challenges and Future Directions

The application of CVAEs for data harmonization in the context of federated learning presents several challenges and offers avenues for future research. While our study demonstrates promising results using synthetic data, further exploration using real-world datasets is necessary to validate the effectiveness of CVAEs in capturing the complex variations present in heterogeneous neuroimaging data.


Furthermore, it is important to investigate the impact of different data characteristics and sources of heterogeneity on the performance of CVAEs in harmonization. This includes



- (a) Principal Component Analysis (PCA) of the feature space of the test dataset after harmonization. The harmonization process effectively aligns the distributions of the features from different institutions, resulting in reduced batch effects.



- (b) PCA of the feature space of the non-harmonized features. Without harmonization, the features exhibit distinct clusters corresponding to different institutions, indicating the presence of significant batch effects.



- (c) PCA of the feature space of the unbiased target features (ground truth). These features represent the ideal scenario where batch effects are completely eliminated, resulting in a homogeneous distribution of data points across different institutions.

Fig. 4.1: Principal Component Analysis (PCA) of the feature space illustrating the effects of harmonization. (a) Shows the harmonized features, (b) shows the non-harmonized features, and (c) shows the unbiased target features. Comparing these plots provides insights into the effectiveness of the harmonization method in reducing batch effects and achieving a more consistent representation of the underlying features.

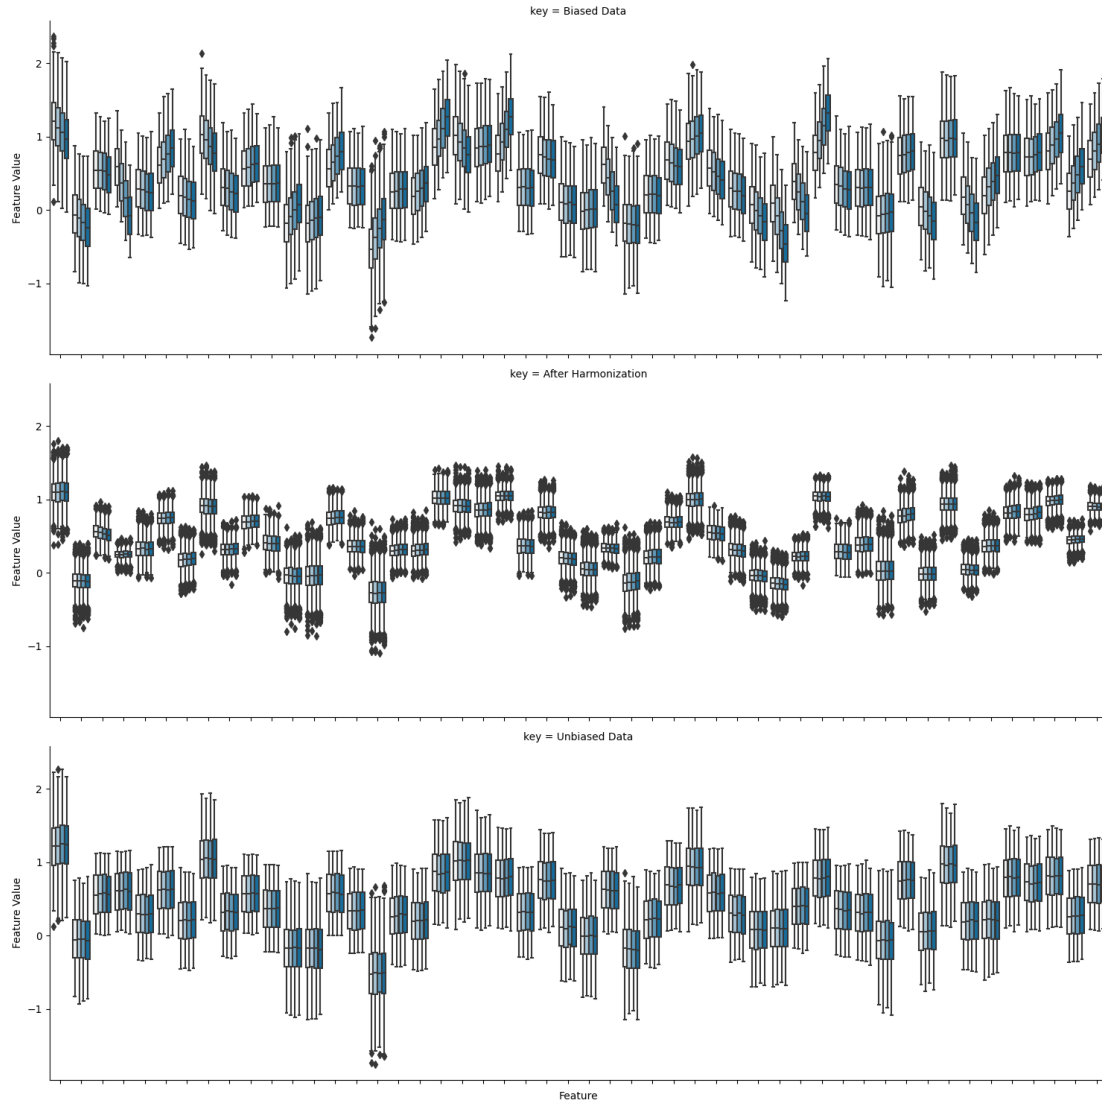
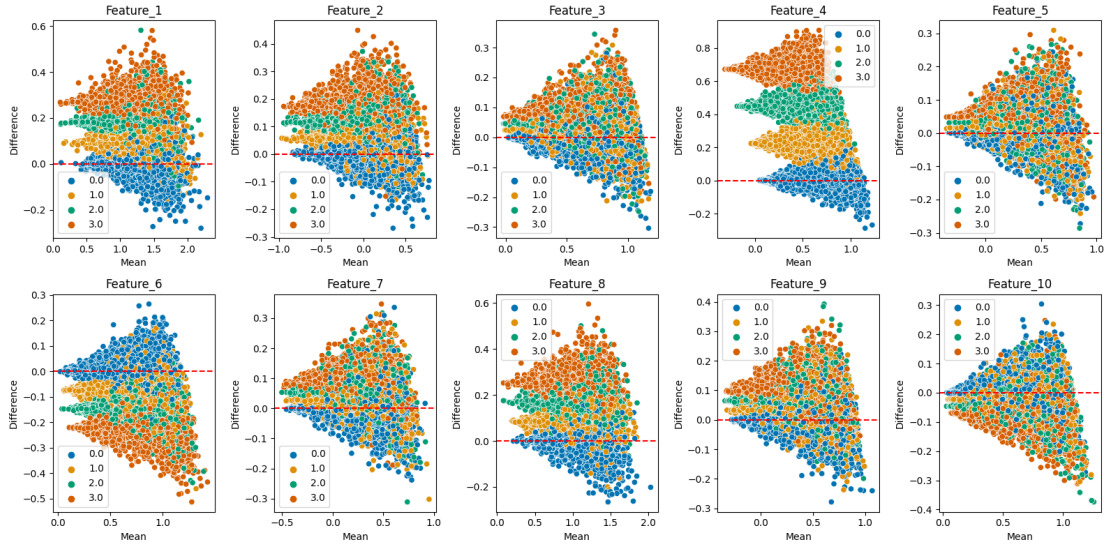
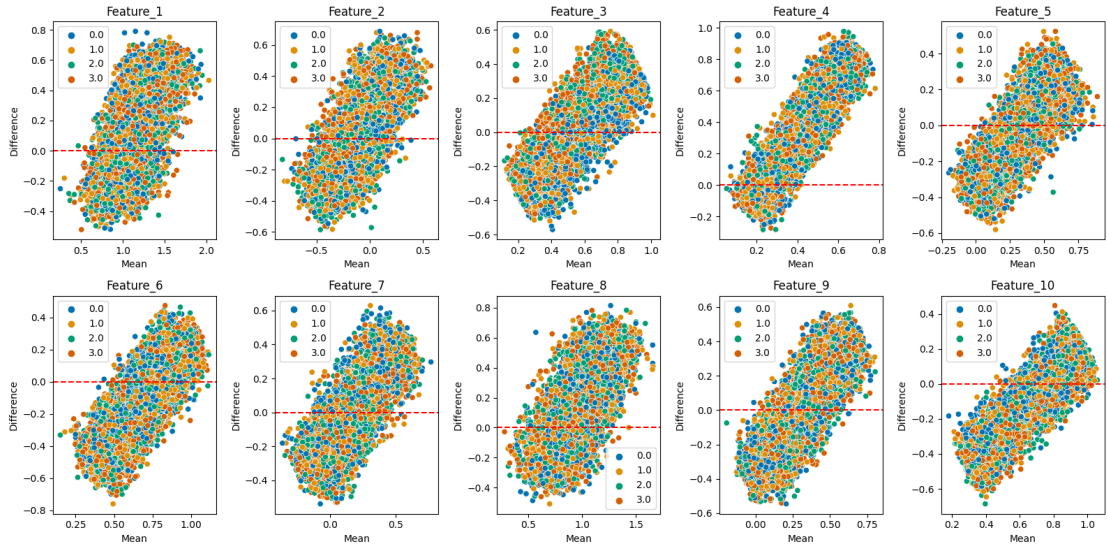


Fig. 4.2: Distribution of features across sites: (top) Before harmonization, (middle) After harmonization, and (bottom) Unbiased distribution. The top plot illustrates the initial feature distribution, where distinct clusters corresponding to different sites are observed. In the middle plot, after applying the harmonization method, the feature distributions are aligned, reducing the variations between sites. The bottom plot represents the ideal scenario with an unbiased distribution, where batch effects have been completely eliminated.



(a) Bland-Altman plot comparing the feature differences between the unbiased dataset and the unharmonized dataset, stratified by batch effects.



(b) Bland-Altman plot comparing the feature differences between the unbiased dataset and the harmonized dataset using the CVAE method, stratified by batch effects.

Fig. 4.3: Bland-Altman plots illustrating the feature differences between the unbiased dataset and the unharmonized dataset (a) and the harmonized dataset (b), stratified by batch effects.

exploring the influence of imbalanced data distributions, missing data, and outliers, as well as understanding how different neuroimaging modalities and preprocessing techniques affect the harmonization process.

In addition to these technical challenges, ethical and privacy considerations are of utmost importance in federated learning approaches. It is crucial to establish robust privacy-preserving mechanisms, data governance frameworks, and regulatory guidelines

to ensure the responsible and secure sharing of sensitive neuroimaging data across institutions while maintaining patient privacy and data protection.

4.6 Conclusions

In this chapter, we have explored the potential of Conditional Variational Autoencoders (CVAEs) for data harmonization in the context of federated learning for neuroimaging. By leveraging the power of CVAEs, we have demonstrated the ability to align and harmonize heterogeneous data from multiple institutions while preserving the underlying features and mitigating the impact of batch effects. Our results on synthetic data have shown promising outcomes, highlighting the potential of CVAEs as a valuable tool for data harmonization.

Data harmonization plays a critical role in biomedical research, enabling the integration and comparison of data from diverse sources. The application of CVAEs in this process offers several advantages. By utilizing the site indicator variable as a control mechanism, CVAEs can effectively capture the variations present in the latent space and align the distributions of the data points from different institutions. This leads to a more consistent representation of the underlying features, enhancing the reproducibility and generalizability of research findings.

While our preliminary exploration has provided insights into the use of CVAEs for data harmonization, there are important challenges and future directions to consider. Real-world validation using large-scale, diverse datasets is necessary to assess the performance and generalizability of CVAEs in capturing the complex variations present in neuroimaging data. Furthermore, addressing computational scalability, exploring the impact of different data characteristics, and addressing ethical and privacy concerns are crucial areas for further research.

In addition to its applications in phenotypical data, it is worth noting that the use of Conditional Variational Autoencoders (CVAEs) for data harmonization in federated learning extends beyond numerical and tabular data. CVAEs can also be effectively applied to image data, leveraging the power of convolutional layers to capture complex spatial patterns and structures.

In conclusion, we explore a potential tool for data harmonization in federated learning, enabling the integration and analysis of multi-centric neuroimaging data while addressing challenges related to batch effects and data heterogeneity. The potential impact of this approach on healthcare outcomes and scientific discoveries is significant, and ongoing

research and collaboration will drive further advancements in the field, leading to improved understanding, diagnosis, and treatment of neurological disorders.

Fed-BioMed: A General Open-source Frontend Framework for Federated Learning in Healthcare

5.1	Introduction	78
5.2	Goals and Contributions of Fed-BioMed	79
5.3	Related works	79
5.4	Implementation Overview and Architecture of Fed-BioMed	80
5.4.1	Client (Node) Service	80
5.4.2	Central Node (Network) Service	81
5.4.3	Researcher Service	81
5.4.4	Communication scheme	82
5.4.5	Security	83
5.4.6	Traceability	83
5.5	Related work	83
5.5.1	Non-medical federated learning frameworks	83
5.5.2	Medical oriented federated learning frameworks	86
5.6	Usage	88
5.7	Applications and Case Studies	88
5.7.1	MNIST analysis	88
5.7.2	Brain imaging data analysis	90
5.8	Results	91
5.9	Conclusions	92
5.10	Supplementary Material	93

Abstract: As the healthcare industry generates data in unprecedented amounts, the feasibility of clinical studies is often impeded by challenges in data access and transfer, particularly in terms of privacy issues. Federated learning provides a privacy-preserving solution for data analysis by employing decentralized optimization methods that keep data securely distributed. Although existing initiatives offer federated learning frameworks, they are often tailored to specific hardware and

modeling approaches and do not natively provide a production-ready environment for deployment.

To address this gap, we propose Fed-BioMed, an open-source federated learning frontend framework specifically designed for healthcare applications. Our framework is built on a versatile architecture that accommodates various models and optimization techniques. We present the software components for client deployment (e.g., healthcare institutions), the central server, and the researcher development environment. We illustrate the workflow for implementing learning models and integrating clients into the federated learning system. Additionally, we showcase a real-world application of our Fed-BioMed framework in the federated analysis of multicentric brain imaging data.

5.1 Introduction

The private and sensitive nature of healthcare information often hampers the use of analysis methods relying on the availability of data in a centralized location. Decentralized learning approaches, such as federated learning, represent today a key working paradigm to empower research while keeping data secure [Li et al., 2020a].

Initially conceived for mobile applications [Konečný et al., 2016], federated learning allows to optimize machine learning models on datasets that are distributed across clients, such as multiple clinical centers in healthcare [Brisimi et al., 2018]. Two main actors play in the federated learning scenario: *clients*, represented for instance by clinical centers, and a *central node* that continuously communicate with the clients [Yang et al., 2019].

Federated learning methods must address three main issues: *security*, by preventing data leakages and respecting privacy policies such as the EU general data protection regulation (GDPR) [Voigt et al., 2017], *communication efficiency*, by optimizing the rounds of communication between clients and the central node, and *heterogeneity robustness*, by properly combining models while avoiding biases from the clients or adversarial attacks aiming to sabotage the models [Bhagoji et al., 2019; Bagdasaryan et al., 2020]. These issues are currently tackled through the definition of novel federated analysis paradigms, and by providing formal guarantees for the associated theoretical properties [Li et al., 2018; Li et al., 2019a].

Besides the vigorous research activity around the theory of federated learning, applications of the federated paradigm to healthcare are emerging [Gupta et al., 2018; Kim et al., 2017; Lee et al., 2018]. Nevertheless, the translation of federated learning to real-life scenarios still requires to face several challenges. Besides the bureaucratic burden, for

federation still requires to establish formal collaboration agreements across partners, the implementation of a federated analysis framework requires to face important technical issues, among which the problem of data harmonization, and the setup of software infrastructures. In particular, from the software standpoint, the practical implementation of federated learning requires the availability of frontend frameworks that can adapt to general modeling approaches and application scenarios, providing researchers with a starting point overcoming problems of deployment, scalability and communication over the internet.

5.2 Goals and Contributions of Fed-BioMed

In this work we propose Fed-BioMed, an open-source production-ready framework for federated learning in healthcare. Fed-BioMed is Python-based and provides modules to deploy general models and optimization methods within federated analysis infrastructures. Besides enabling standard federated learning aggregation schemes, such as federated averaging (FedAVG) [Konečný et al., 2016; McMahan et al., 2017], Fed-BioMed allows the integration of new models and optimization approaches. It is also designed to enable the integration with currently available federated learning frameworks, while guaranteeing secure protocols for broadcasting.

We expect this framework to foster the application of federated learning to real-life analysis scenarios, easing the process of data access, and opening the door to the deployment of new approaches for modeling and federated optimization in healthcare. The code will be freely accessible from our repository page (<https://gitlab.inria.fr/fedbiomed>).

5.3 Related works

NVIDIA Clara is a large initiative focusing on the deployment of federated learning in healthcare [Yuhong et al., 2019], currently providing a service where users can deploy personalized models. The code of the project is not open, and it requires the use of specific hardware components. This may reduce the applicability of federated learning to general use-cases, where client's facilities may face restrictions in the use of proprietary technology.

The Collaborative Informatics and Neuroimaging Suite Toolkit for Anonymous Computation (COINSTAC) [Plis et al., 2016] focuses on single-shot and iterative optimization schemes for decentralized multivariate models, including regression and latent variable analyses (e.g. principal/independent-component analysis, PCA-ICA, or canonical

correlation analysis, CCA). This project essentially relies on distributed gradient-based optimization as aggregating paradigm. A frontend distributed analysis framework for multivariate data modeling was also proposed by [Silva et al., 2019]. Similarly to COINSTAC, the framework focuses on the federation of multivariate analysis and dimensionality reduction methods. The PySyft initiative [Ryffel et al., 2018] provides an open-source wrapper enabling federated learning as well as encryption and differential privacy. At this moment however this framework focuses essentially on optimization schemes based on stochastic gradient descent, and does not provide natively a deployable production-ready framework. Fed-BioMed is complementary to this initiative and allows interoperability, for example by enabling unit testing based on PySift modules.

5.4 Implementation Overview and Architecture of Fed-BioMed

Fed-BioMed is designed as a microservices-based architecture, organized according to the different actors in Federated Learning (FL): (i) the **clients (nodes)**, responsible for managing the data they want to allow modeling on, (ii) a server or **central node (network) service**, in charge of managing communication and model parameter sharing, and (iii) the **researcher**, whose model or optimization method is to be deployed and tested. The architecture is inspired by the 2018 “Guide for Architectural Framework and Application of Federated Machine Learning” (Federated Learning infrastructure and Application standard)¹, basing it on a server-client paradigm with three types of instances: central node, clients, and federators.

5.4.1 Client (Node) Service

Clients, nodes or institutes are responsible for storing the datasets through a dedicated client application. Since not every client is required to store the same type of data, this enables studies with heterogeneous or missing features across participants. Moreover, depending on the target data source and on the analysis paradigm, centers can be either included or ignored from the study. Each client verifies the number of current jobs in queue with the central node, as well as data types and models associated to the running jobs.

Figure 5.2 shows Fed-BioMed’s user interface which allows institutions to configure multiple types of data for modeling consent, such as plain-text based data in the form of CSV or TSV files, default datasets from third-party repositories, or image data. Once configured, a client can start and stop the service whenever considered appropriate.

¹https://standards.ieee.org/project/3652_1.html

Clients receive commands from researchers that trigger defined events like: ping to check node availability, local training of a specific model, and data description requests (e.g., sample sizes to account for size effects).

5.4.2 Central Node (Network) Service

The central node or network service serves as a trusted broker between researchers and participating clients or institutions. It manages communication channels, model updates, and serves as an ephemeral storage service for model aggregation, ensuring secure and efficient collaboration between all parties involved.

In order to maintain an optimized communication channel, event triggering, messaging, and model parameter sharing are managed by two services. Both commands and data flow can coexist without hindering each other. Messaging is managed through MQTT, a standard industrial protocol for efficient messaging communication. Additionally, model flow is managed using HTTP requests via a REST API developed in Django as the backend framework. This architecture consequently delegates parameters' data flow to the central node instead of being the researcher's responsibility.

These services are also provided within Fed-BioMed in the form of Docker containers, so they can be scaled horizontally. The central node's role in a federated learning setup ensures secure and efficient communication between all parties involved in the study.

Furthermore, the central node is responsible for maintaining a job queue, which keeps track of the current jobs in progress and their associated models and data. This allows clients to verify the number of current jobs in the queue, as well as data types and models related to the running jobs.

5.4.3 Researcher Service

The researcher service in Fed-BioMed provides a comprehensive and efficient platform for researchers to develop, deploy, and manage their models in a federated learning environment. By providing support for various aggregation methods, model development, and orchestration of the training process, Fed-BioMed enables researchers to effectively address data heterogeneity, ensure robust model performance across diverse nodes, and maintain control over the entire federated learning process.

An `Experiment` class is included in the package and it is responsible for orchestrating the training process on available nodes. Such orchestration involves:

- Searching for datasets on active nodes based on specific tags provided by a researcher and used by the nodes to identify the dataset.
- Uploading the training plan file created by the researcher and sending the file URL to the nodes.
- Sending model and training arguments to the nodes.
- Tracking the training process in the nodes during all training rounds.
- Verifying the nodes' responses to ensure that each round is successfully completed in every node.
- Downloading the local model parameters after every round of training.
- Aggregating the local model parameters based on the specified federated approach, and eventually sending the aggregated parameters to the selected nodes for the next round.

As of the publication of this work, Fed-BioMed supports several aggregation methods to tackle the challenges posed by heterogeneous data in federated learning. The first method, FedAVG, is a widely-used standard aggregation scheme in federated learning that performs federated averaging. Additionally, we provide FedProx and SCAFFOLD aggregation methods to further enhance the capabilities of the researcher service in handling diverse data distributions and sample sizes.

5.4.4 Communication scheme

Every instance is packaged and deployed in form of Docker containers [Merkel, 2014] interacting between each other through HTTP requests. Containerized instances help to overcome software/hardware heterogeneity issues by creating an isolated virtualized environment with a predefined operating system (OS), thus improving reproducibility as every center run on the same software environment. This scheme also achieves scalability and modularity when dealing with large amounts of clients. In this case, multiple instances of the API can be created under a load balancer, while federator instances can be separately deployed on a dedicated computation infrastructure.

5.4.5 Security

Fed-BioMed addresses typical security issues regarding communication (e.g. man-in-the-middle and impersonation attacks) and access permissions as follows: 1) all requests are encrypted by using HTTP over SSL, 2) user authentication relies on a password-based scheme, and 3) reading/writing operations are restricted by role definitions. Protection from adversarial or poisoning attacks [Bhagoji et al., 2019] is currently not in the scope of this work, but can be naturally integrated as part of the federator. In the future, malicious attacks will be also prevented by implementing certification protocols attesting the safety of the model source code before deployment [Shen et al., 2016].

5.4.6 Traceability

To allow transparency to the centers and for the sake of technical support, each instance leverages on a logging system that allows to keep track of every request made, shared data and of the current available jobs.

The architecture behind Fed-BioMed is illustrated in Figure 5.1. The common procedure involves the deployment of one or multiple jobs from the researchers. Each job must contain the architecture model to be trained and its initialized parameters for reproducibility, the number of rounds, and the federator instance to be used as optimizer.

5.5 Related work

The interest in federated learning (FL) has led to a significant increase in the number of FL frameworks in recent years. Some well-established frameworks, such as TensorFlow Federated (TFF)², FedML [He et al., 2020], IBM-FL [Ludwig et al., 2020], FATE [Liu et al., 2021], PaddleFL [Ma et al., 2019], and PySyft [Ziller et al., 2021], are not specifically tailored for biomedical applications but could potentially be deployed in this domain. Alternatively, more recent works that were not considered at the time of this work have emerged and have been applied to medical imaging such as SubstraFL [Galtier et al., 2019], OpenFL [Reina et al., 2021], and Flare [Roth et al., 2022].

5.5.1 Non-medical federated learning frameworks

IBM Federated Learning

IBM-FL [Ludwig et al., 2020] is a Python framework designed for federated learning

²<https://www.tensorflow.org/federated>




Fig. 5.1: Clients providing different data sources share their local model parameters with the central node. The central node creates the jobs that will be run by the clients, and transmits the initialization parameters for the models in training. The federator gathers the collected parameters and combines them into a global model that is subsequently shared with the clients for the next training round. As instances are isolated in containers, new instances, such as a new federators (dashed line), can be introduced without altering the behavior of the infrastructure.

in enterprise environments. The framework provides a basic fabric for FL, upon which advanced features can be built. It is not dependent on any specific machine learning framework and supports various learning topologies, such as a shared aggregator, and protocols. IBM Federated Learning aims to offer a solid foundation for federated learning, enabling a wide variety of learning models, topologies, and learning models, particularly in enterprise and hybrid-cloud settings.

PySyft

PySyft[Ziller et al., 2021] is a popular choice, serving as a wrapper for TensorFlow and PyTorch in federated learning scenarios. While PySyft offers a simple wrapper, it lacks security and scalability features for production environments due to its reliance on a single channel for parameter transmission and event triggering. Furthermore, it only supports multiple nodes in development mode through Virtual workers. PySyft's production solution, Duet, permits only a single node-to-researcher connection, which is suboptimal for FL in production settings.

TensorFlow Federated

TensorFlow Federated (TFF) is an open-source framework designed for machine learning and other computations on decentralized data. It aims to facilitate research and experimentation with Federated Learning (FL), an approach where a shared global model is

Data Management

Here you can manage the data associated to your center. It is important to take the following into consideration:

- Data must be formatted in CSV (`.csv`) or compressed CSV (`.csv.gz`)
- Subjects ID must be in the first column of the CSV
- If you provide a label/diagnosis column, mark it with the prefix `DX` or `label`

Your data will not be shared outside this facility. Instead, this will generate a local backup to work on.

Select the data type you want to configure/update

Brain Data

Data type: Features

Insert the name of the data file (or part of it) and press ENTER. We will look for all of them in the data path.

corr

These are the files found. Please select the one belonging to ADNI and OASIS:

data _corrected_std.csv (1.5 MB)

Update my data

Fig. 5.2: A screenshot of the data management utility that is integrated within the Fed-BioMed framework. This utility plays a crucial role in facilitating the secure and efficient transfer of data between different healthcare institutions and research centers. With the Fed-BioMed data management utility, users can securely upload, access, and share data with other users within a federated learning setup. The utility also includes features such as data pre-processing, quality control, and privacy-preserving mechanisms to ensure that sensitive information is protected at all times.

trained across multiple participating clients while retaining their training data locally. TFF allows developers to simulate federated learning algorithms on their models and data and experiment with novel algorithms.

The TFF platform comprises two layers: the Federated Learning (FL) layer, which offers high-level interfaces for integrating existing Keras or non-Keras machine learning models into the TFF framework, and the Federated Core (FC) layer, providing lower-

level interfaces for expressing custom federated algorithms. The FC layer combines TensorFlow with distributed communication operators within a strongly-typed functional programming environment.

Currently, TFF is designed for experimentation purposes and does not include tools for production settings, such as deployment on mobile phones. The framework supports a local simulation runtime for experiments. It is expected that the open-source ecosystem surrounding TFF will evolve to incorporate runtimes targeting physical deployment platforms.

Flower

Flower[Beutel et al., 2020], another FL framework, utilizes the same communication channel for both event triggering and model parameter transmission. Although Flower potentially supports a large number of nodes, researchers must distribute the same script across nodes to run an experiment, complicating deployment. Additionally, Flower requires file paths to be specified or hard-coded at each node for data loading, making it unsuitable for continuous delivery-centric applications. The framework also delegates messaging service management to the researcher, reducing system resiliency in the event of researcher connection or availability issues.

FATE (Federated AI Technology Enabler)

FATE[Liu et al., 2021] offers a production-grade package enabling the deployment of existing multi-party computation (MPC) and FL solutions, including homomorphic encryption capabilities. However, its focus on predefined methods limits flexibility for research purposes.

Paddle FL

Paddle FL[Ma et al., 2019] presents an industrial FL solution specifically designed for the PaddlePaddle³ machine learning library. While it provides features such as support for differential privacy, Paddle FL remains exclusive to the PaddlePaddle framework, lacking compatibility with PyTorch or TensorFlow.

5.5.2 Medical oriented federated learning frameworks

SubstraFL

SubstraFL [Galtier et al., 2019], developed by Owkin, is a Python library based on the Substra library. Currently SubstraFL is utilized in healthcare applications such as drug discovery within the MELLODY project and oncology. SubstraFL's architecture is built upon collaboration, privacy, and traceability. While it does not provide any tools

³<https://github.com/PaddlePaddle/Paddle>

	FED-BIOMED	Flower	Tensorflow FL	FATE	PaddleFL	PySyft
Federated Learning:						
Gradient-descent based models	●	●	●	●	●	●
Support for multiple ML libraries	●	●				●
Node sampling strategies support	●	●	●			
Aggregation methods support	●	●	●			●
Model agnostic	●					●
Privacy:						
<i>f</i> -DP mechanisms support, MPC	●				●	●
Other features:						
Ready for production (deployable)	●			●	●	●
Node data consent management	●					●
Non-gradient-descent based models	●	●				
VPN support provided	●					

Tab. 5.1: Comparison of Existing Frameworks in Federated Learning (FL) as of July 2021. Green circles represent included features, while yellow circles indicate expected features. Flare was not considered due to its closed-source nature and limited compatibility with NVIDIA hardware.

dedicated to healthcare or biomedical assets, its distributed ledger-based architecture ensures secure and transparent operations. However, its centralized governance and roadmap may hinder the development of an open-source community.

OpenFL

OpenFL [Reina et al., 2021] is a Python-based FL library, was originally designed for healthcare applications and later expanded to be more general from the use case. OpenFL's strong focus on cybersecurity, achieved through measures like TLS-encrypted communication and PKI certificates which is optimal for multi-institutional collaborations. While tutorials are provided for medical applications, no core functionalities are specific to biomedical data sources. Additionally, the adoption of OpenFL may be limited due to its reliance on proprietary hardware.

Flare

Flare, another Python-based FL library developed by NVIDIA, has been used in various healthcare applications, including the development of a triaging model for COVID-19 patients [Dayan et al., 2021] and classification and segmentation tasks on medical images [Roth et al., 2020; Sarma et al., 2021]. Flare's design principles focus on scalability, flexibility, and lightweight, making it suitable for cross-silo FL in production settings and simulated FL for researchers. Although Flare has been showcased in healthcare settings, it is still a generic framework without specific tools for biomedical data sources. Furthermore, Flare's reliance on specialized hardware and configuration files may limit its flexibility and applicability in hospital settings.

5.6 Usage

Deploying a new model.

Fed-BioMed relies on a common convention for deploying a new model. A model must be defined as a PyTorch [Paszke et al., 2019] module class, containing the common methods for any torch module:

- `__init__(**kwargs)` method: defines the model initialization parameters;
- `forward()` method: provides instructions for computing local model updates.

Once the new class is defined, it can be integrated in the model zoo for both clients and federators (Figure 5.3, left).

Deploying a new federator or aggregation function in the backend is obtained by creating a containerized service that queries the API for the necessary submissions at each round, and subsequently aggregates the submitted local updates (Figure 5.3, right).

5.7 Applications and Case Studies

This section showcases Fed-BioMed through two different experiments. The first experiment involves analyzing the MNIST dataset [LeCun et al., 2010] with the participation of 25 clients. The second experiment involves a multi-centric analysis of brain imaging data, conducted in collaboration with four research partners located in different geographical regions."

5.7.1 MNIST analysis

The 60000 MNIST images were equally split among 25 centers. Each center was synthetically emulated and setup in order to interact with the centralized API. The model was represented by a variational autoencoder (VAE) implementing non-linear encoding and decoding functions composed by three layers respectively, with a 2-dimensional associated latent space. For training we used a learning rate of $l = 1 \times 10^{-3}$, 10 local epochs for 30 optimization rounds, while the federated aggregating scheme was FedAVG.

```

class MyModel(torch.nn.Module):
    def __init__(**kwargs):
        # Define model

    @staticmethod
    def loss_function(**args):
        # Define a personalized
        # loss function if needed

    def forward(self, **kwargs):
        # Compute the local updates:
        # parameters to aggregate

    # You can also define
    # other functions
    # (e.g. encode, decode)

```

```

from fed_requests import fedbionet_api as api

# Get list of running Jobs
jobs = api.get_jobs(
    params={
        'status': 'Running',
        'federator': 'MY_METHOD'
    }
)
for job in jobs:
    # Get submissions made by clients
    job_url = job['url']
    current_round = job['current_round']
    submissions = api.get_submissions(
        job_url=job_url,
        current_round=current_round
    )

    # Aggregate submitted models

```

Fig. 5.3: Examples of model declaration (left) and creation of a new federator (right).

	Center 1	Center 2	Center 3	Center 4
No. of participants (M/F)	448/353	454/362	1070/930	573/780
Clinical status				
No. healthy	175	816	2000	695
No. MCI and AD	621	0	0	358
Age \pm sd (range) [years]	73.74 \pm 7.23	28.72 \pm 3.70	63.93 \pm 7.49	67.58 \pm 10.04
Age range [years]	54 - 91	22 - 37	47 - 81	43 - 97

Tab. 5.2: Demographic information for each of the centers that participated in training models using their MRI-derived brain data is provided below. The cohorts include individuals with Mild Cognitive Impairment (MCI) and Alzheimer’s Disease (AD).

5.7.2 Brain imaging data analysis

In this experiment we use our framework to perform dimensionality reduction in multi-centric structural brain imaging data (MRI) across datasets from different geographical locations, providing cohorts of healthy individuals and patients affected by cognitive impairment.

Four centers participated to the study and were based geographically as follows: two centers in France (Center 1 and Center 4), one in the UK (Center 3), and a last one in the US (Center 2). The central node and the federator were also located in France. Each center was running on different OS and the clients were not 100% online during the day allowing to test the robustness of the framework in resuming the optimization in real-life conditions.

For each center, data use permission was obtained through formal data use agreements. Data characteristics across clients are reported in Table 5.2. A total of 4670 participants were part of this study, and we note that the data distribution is heterogeneous with respect to age, range and clinical status. 92 features subcortical volumes and cortical thickness were computed using FreeSurfer [Fischl, 2012] and linearly corrected by sex, age and intra-cranial volume (ICV) at each center.

This analysis involved data standardization with respect to the global mean and standard deviation computed across centers, and dimensionality reduction was performed via a VAE implementing a linear embedding into a 5-dimensional latent space. Federated data standardization was implemented as in [Silva et al., 2019], while VAE’s parameters aggregation was performed through FedAVG. Federated learning was performed by specifying a pre-defined budget of 30 rounds of client-server iterations in total with 15 epochs/client-round at a learning rate of $l = 1 \times 10^{-3}$.

5.8 Results

The evolution of the models during training can be assessed by analyzing the weights' norm across iterations (Figure 5.4 and supplementary material for the complete set of weights). MNIST parameters evolution is shown in the top panel, while the related test set projected onto the latent space is shown in Figure 5.5, left panel, describing a meaningful variability across digits and samples.

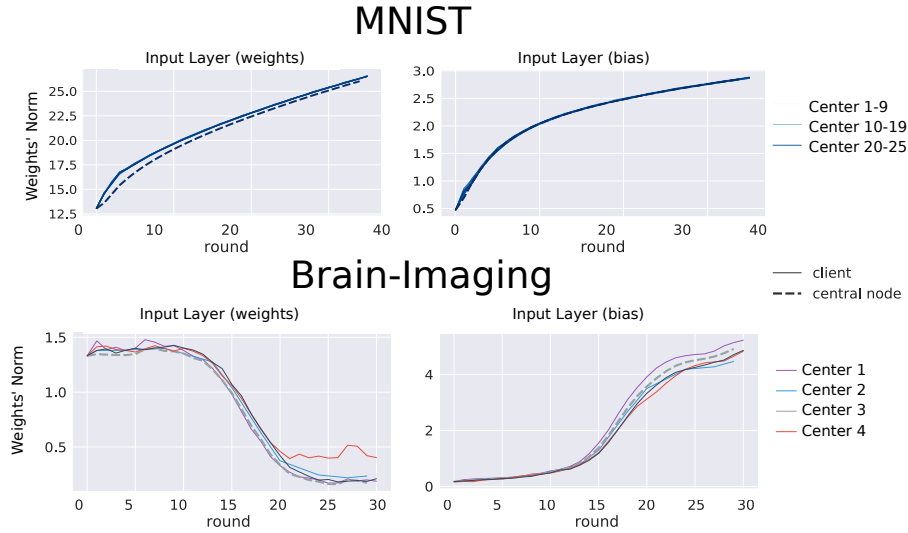


Fig. 5.4: Illustration of parameter evolution for VAE parameters (input layer). The federated model closely follows the clients' weights distribution. **Top:** MNIST across 25 centers with equally distributed data. **Bottom:** brain-imaging heterogeneous dataset across 4 centers with unevenly distributed data. Continuous lines: clients weights' norm. Dashed lines: federated model weights' norm.

Concerning the brain imaging data analysis experiment, the evolution of the encoding parameters throughout the 30 optimization rounds is shown in Figure 5.4, bottom panel. For this real-world application we also collected each client's elapsed time to report its local updates to the central node, as well as the average time per round in the best scenario (Figure 5.6). As expected, the clients' time varies depending on the geographical proximity with the central node, as well as on the local upload/download speed. The model was further investigated by inspecting the latent space on the subset of the training data available to the coordinating Center 1. The right panel of Figure 5.5 shows that although most of the training data for the VAE comes from healthy and young participants, the model is also able to capture the pathological variability related to cognitive impairment in aging. The latent variables associated to the observations of Center 1 indeed show significantly different distributions across different clinical groups, from healthy controls (CN), to subjects with mild cognitive impairment (MCI), and patients with Alzheimer's disease (AD).

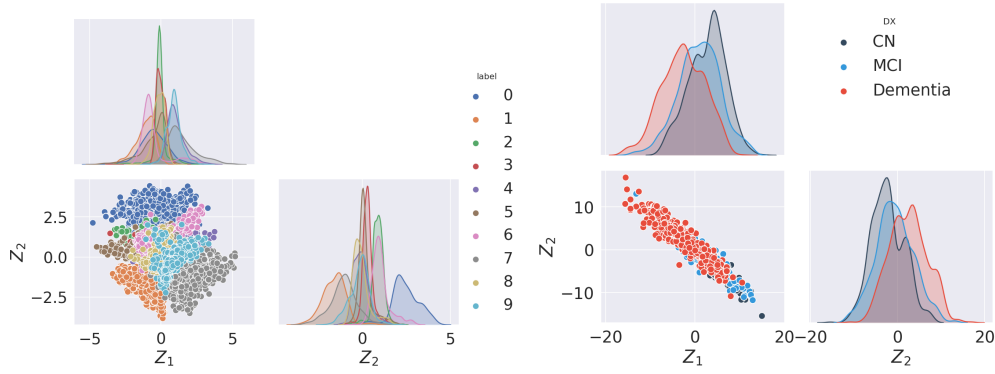


Fig. 5.5: **Left:** MNIST pixel data projected onto the latent space. **Right:** Brain features of Center 1 projected onto the first 2 components in the latent space. Although the model was trained with unbalanced data, it is still able to capture pathological variability. CN: healthy controls; MCI: mild cognitive impairment; Dementia: dementia due to with Alzheimer’s disease.




Fig. 5.6: **Top:** User average elapsed time per round (since a new version of the model is made available and each user to submit its local update). Geographical and data distribution: Center 1 (FR) , Center 2 (US), Center 3 (UK) and Center 4 (FR). **Bottom:** Averaged elapsed time across centers per round of updates.

5.9 Conclusions

This work presents an open-source framework for deploying general federated learning studies in healthcare, providing a production ready reference to deploy new studies based on federated models and optimization algorithms. Our experimental results show that the framework is stable in communication, while being robust to handle clients going temporarily out of the grid. Scalability is obtained thanks to the use of containerized services. The ability to handle client-authentication and the use of secured broadcasting protocols are also appealing security features of Fed-BioMed. While the experiments mostly focused on VAE and Federated Averaging as aggregating paradigm, our framework is completely extensible to other distributed optimization approaches.

Future work will therefore integrate additional models for the analysis of different data modalities and bias, as well as enhanced secure P2P encryption. Concerning the clinical experimental setup, the brain application was chosen to provide a demonstration of our framework to a real-life analysis scenario, and it is not aimed to address a specific clinical question. In the future, the proposed work Fed-BioMed will be a key component for clinical studies tailored to address challenging research questions, such as the analysis of imaging-genetics relationships in current meta-analysis initiatives [Thompson et al., 2014].

5.10 Supplementary Material

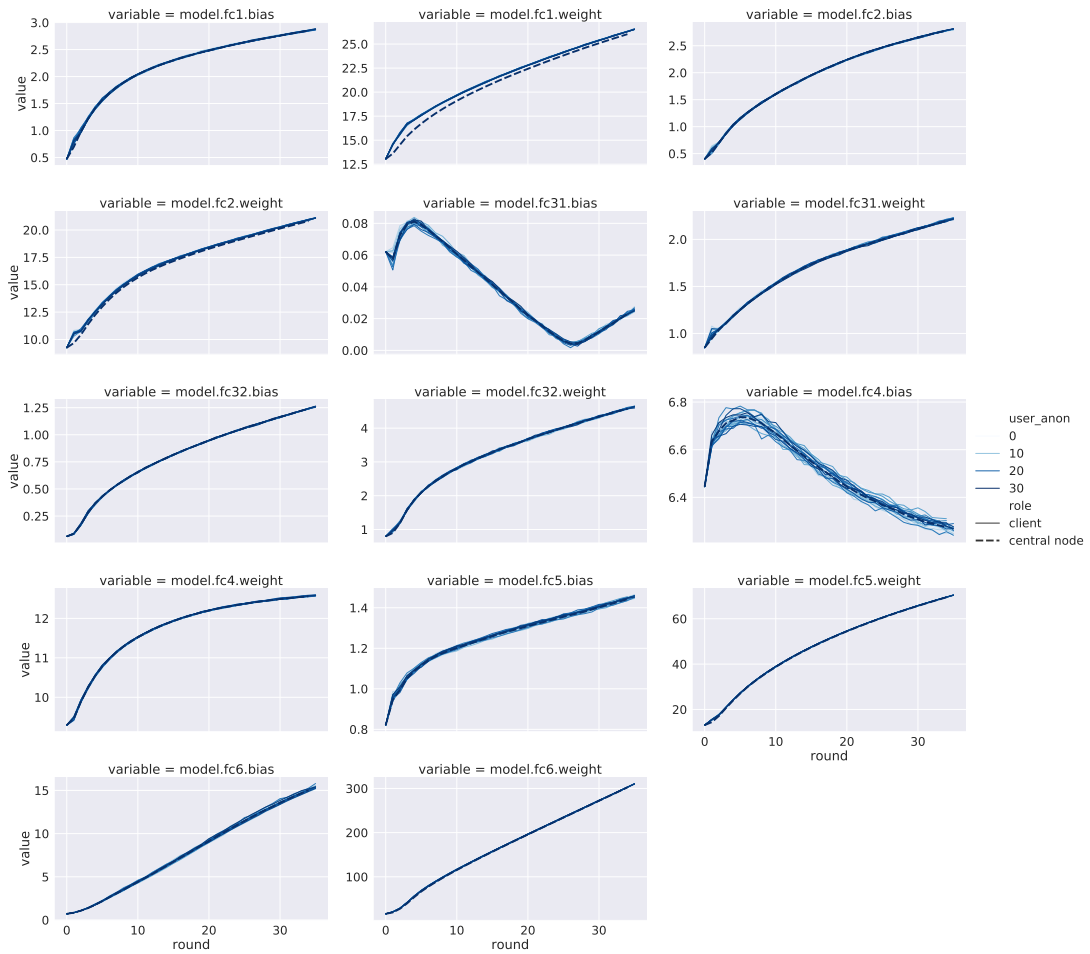


Fig. 5.7: Parameter convergence for encoding-decoding VAE parameters. The federated model closely follows the clients' weights distribution. MNIST across 25 centers with equally distributed data. Dashed lines: federated model weights' norm.

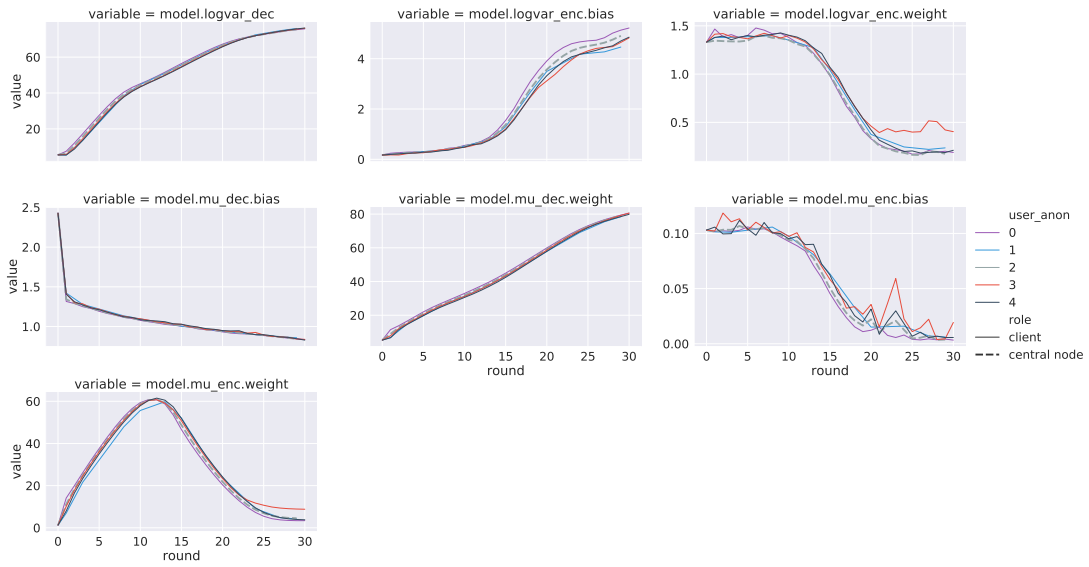


Fig. 5.8: Parameter convergence for encoding-decoding VAE parameters. Brain-imaging heterogeneous dataset across 4 centers with unevenly distributed data. Continuous lines: clients weights' norm. Dashed lines: federated model weights' norm.

Conclusion

In this thesis, we have presented a comprehensive investigation of federated learning and data harmonization techniques for multi-centric neuroimaging studies. By addressing the challenges associated with data heterogeneity, batch effects, and privacy concerns, our research has contributed to the development of more efficient, robust, and scalable approaches for the analysis of large-scale neuroimaging data.

Summary of the Main Contributions

Throughout the thesis, we have covered the following key contributions:

- Introduced two methods for addressing data heterogeneity in federated learning setups, Fed-ComBat and federated mixed-effect modeling, which aim to perform data harmonization and modeling on heterogeneous data respectively.
- Presented a black-box optimization scheme for federated learning, based on gradient-free optimization of a global model, to improve the optimization process in federated setups, enabling institutions and hospitals with limited computational resources to participate in federated learning while achieving accurate and generalizable models.
- Developed Fed-BioMed, an open-source framework for federated learning in health-care, which fulfills the need for a common collaborative framework that is compliant with privacy and ethical standards, empowering real-life federated applications.

Perspectives and Future Applications

Integration of methods into the Fed-BioMed ecosystem

One potential direction for future work is the seamless integration of the developed methods into the Fed-BioMed ecosystem. This would enable researchers to apply these methods using a standardized framework, facilitating collaboration and fostering the adoption of federated learning in healthcare research. By incorporating Fed-ComBat, federated mixed-effect modeling, and the black-box optimization scheme into Fed-BioMed, researchers can more effectively address data heterogeneity, optimize model performance, and overcome computational resource limitations.

Expanding the Impact on Large Multicentric Consortia

Another promising avenue for future research is the application of the methods presented in this thesis to large multicentric consortia such as the Enhancing NeuroImaging Genetics through Meta-Analysis (ENIGMA) consortium. By leveraging federated learning techniques and the Fed-BioMed framework, researchers within these consortia can work collaboratively on large-scale neuroimaging projects while preserving data privacy and maintaining compliance with data protection regulations. The methods proposed in this thesis can help consortia like ENIGMA to tackle data heterogeneity, optimize models, and share computational costs, ultimately leading to more efficient and effective advancements in healthcare.

Overall, this thesis has provided methodological and technical contributions that address the challenges of data heterogeneity, optimization, and infrastructure in federated learning setups for neuroimaging research. Our work aims to facilitate more efficient and effective advancements in healthcare while preserving patient privacy and data governance. By developing novel methods and frameworks, and integrating them into real-life applications and large multicentric consortia, we hope to pave the way for the broader adoption of federated learning in neuroimaging research, ultimately leading to improved understanding, diagnosis, and treatment of neurological disorders.

Bibliography

- [Ashburner et al., 2005] John Ashburner and Karl J Friston. “Unified segmentation”. In: *Neuroimage* 26.3 (2005), pp. 839–851 (cit. on p. 32).
- [Auer, 2002] Peter Auer. “Using confidence bounds for exploitation-exploration trade-offs”. In: *Journal of Machine Learning Research* 3.Nov (2002), pp. 397–422 (cit. on p. 55).
- [Bagdasaryan et al., 2020] Eugene Bagdasaryan et al. “How to backdoor federated learning”. In: *International Conference on Artificial Intelligence and Statistics*. 2020, pp. 2938–2948 (cit. on p. 78).
- [Bajaj et al., 2021] Ishan Bajaj, Akhil Arora, and MM Faruque Hasan. “Black-box optimization: Methods and applications”. In: *Black box optimization, machine learning, and no-free lunch theorems*. Springer, 2021, pp. 35–65 (cit. on p. 51).
- [Bergstra et al., 2012] James Bergstra and Yoshua Bengio. “Random search for hyperparameter optimization.” In: *Journal of machine learning research* 13.2 (2012) (cit. on pp. 51, 52).
- [Bethlehem et al., 2022] Richard AI Bethlehem et al. “Brain charts for the human lifespan”. In: *Nature* 604.7906 (2022), pp. 525–533 (cit. on pp. 23, 41).
- [Beutel et al., 2020] Daniel J Beutel et al. “Flower: A friendly federated learning research framework”. In: *arXiv preprint arXiv:2007.14390* (2020) (cit. on p. 86).
- [Bhagoji et al., 2019] Arjun Nitin Bhagoji et al. “Analyzing federated learning through an adversarial lens”. In: *International Conference on Machine Learning*. 2019, pp. 634–643 (cit. on pp. 78, 83).
- [Bonawitz et al., 2017] Keith Bonawitz et al. “Practical secure aggregation for privacy-preserving machine learning”. In: *proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*. 2017, pp. 1175–1191 (cit. on p. 17).
- [Brisimi et al., 2018] Theodora S Brisimi et al. “Federated learning of predictive models from federated electronic health records”. In: *International journal of medical informatics* 112 (2018), pp. 59–67 (cit. on pp. 9, 15, 78).

- [Brochu et al., 2010] Eric Brochu, Vlad M Cora, and Nando De Freitas. “A tutorial on Bayesian optimization of expensive cost functions, with application to active user modeling and hierarchical reinforcement learning”. In: *arXiv preprint arXiv:1012.2599* (2010) (cit. on p. 52).
- [Bukaty, 2019] Preston Bukaty. *The California Consumer Privacy Act (CCPA): An implementation guide*. IT Governance Publishing, 2019. URL: <http://www.jstor.org/stable/j.ctvjghvnn> (visited on Oct. 31, 2022) (cit. on p. 22).
- [Bull, 2011] Adam D Bull. “Convergence rates of efficient global optimization algorithms.” In: *Journal of Machine Learning Research* 12.10 (2011) (cit. on p. 50).
- [Calzada, 2022] Igor Calzada. “Citizens’ data privacy in China: The state of the art of the Personal Information Protection Law (PIPL)”. In: *Smart Cities* 5.3 (2022), pp. 1129–1150 (cit. on p. 22).
- [Chang et al., 2018] Ken Chang et al. “Distributed deep learning networks among institutions for medical imaging”. In: *Journal of the American Medical Informatics Association* 25.8 (2018), pp. 945–954 (cit. on p. 15).
- [Chen et al., 2022] Andrew A Chen et al. “Privacy-preserving harmonization via distributed ComBat”. In: *Neuroimage* 248 (2022), p. 118822 (cit. on pp. 23, 24, 27).
- [Cheung et al., 2023] Alexander TM Cheung et al. “Methods and impact for using federated learning to collaborate on clinical research”. In: *Neurosurgery* 92.2 (2023), pp. 431–438 (cit. on p. 23).
- [Costafreda, 2009] Sergi G Costafreda. “Pooling FMRI data: meta-analysis, mega-analysis and multi-center studies”. In: *Frontiers in neuroinformatics* (2009), p. 33 (cit. on p. 9).
- [Dale et al., 1999] Anders M Dale, Bruce Fischl, and Martin I Sereno. “Cortical surface-based analysis: I. Segmentation and surface reconstruction”. In: *Neuroimage* 9.2 (1999), pp. 179–194 (cit. on p. 32).
- [Dayan et al., 2021] Ittai Dayan et al. “Federated learning for predicting clinical outcomes in patients with COVID-19”. In: *Nature medicine* 27.10 (2021), pp. 1735–1743 (cit. on p. 87).
- [Desikan et al., 2006] Rahul S Desikan et al. “An automated labeling system for subdividing the human cerebral cortex on MRI scans into gyral based regions of interest”. In: *Neuroimage* 31.3 (2006), pp. 968–980 (cit. on p. 32).
- [Di Martino et al., 2014] Adriana Di Martino et al. “The autism brain imaging data exchange: towards a large-scale evaluation of the intrinsic brain architecture in autism”. In: *Molecular psychiatry* 19.6 (2014), pp. 659–667 (cit. on p. 32).

- [Dwork, 2006] Cynthia Dwork. “Differential privacy”. In: *Automata, Languages and Programming: 33rd International Colloquium, ICALP 2006, Venice, Italy, July 10-14, 2006, Proceedings, Part II* 33. Springer. 2006, pp. 1–12 (cit. on p. 10).
- [Eisenhauer, 2021] Joseph G Eisenhauer. “Meta-analysis and mega-analysis: A simple introduction”. In: *Teaching Statistics* 43.1 (2021), pp. 21–27 (cit. on p. 11).
- [El Emam et al., 2014] Khaled El Emam and Bradley Malin. “Concepts and methods for de-identifying clinical trial data”. In: *Paper commissioned by the Committee on Strategies for Responsible Sharing of Clinical Trial Data* (2014) (cit. on p. 9).
- [Ellis et al., 2009] Kathryn A Ellis et al. “The Australian Imaging, Biomarkers and Lifestyle (AIBL) study of aging: methodology and baseline characteristics of 1112 individuals recruited for a longitudinal study of Alzheimer’s disease”. In: *International psychogeriatrics* 21.4 (2009), pp. 672–687 (cit. on p. 32).
- [European Commission, 2016] European Commission. *Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation) (Text with EEA relevance)*. 2016. URL: <https://eur-lex.europa.eu/eli/reg/2016/679/oj> (cit. on p. 22).
- [Fischl, 2012] Bruce Fischl. “FreeSurfer”. In: *Neuroimage* 62.2 (2012), pp. 774–781 (cit. on p. 90).
- [Fischl et al., 2004] Bruce Fischl et al. “Automatically parcellating the human cerebral cortex”. In: *Cerebral cortex* 14.1 (2004), pp. 11–22 (cit. on p. 32).
- [Fortin et al., 2014] Jean-Philippe Fortin et al. “Functional normalization of 450k methylation array data improves replication in large cancer studies”. In: *Genome biology* 15.11 (2014), pp. 1–17 (cit. on p. 67).
- [Fortin et al., 2017] Jean-Philippe Fortin et al. “Harmonization of multi-site diffusion tensor imaging data”. In: *Neuroimage* 161 (2017), pp. 149–170 (cit. on p. 23).
- [Fortin et al., 2018] Jean-Philippe Fortin et al. “Harmonization of cortical thickness measurements across scanners and sites”. In: *Neuroimage* 167 (2018), pp. 104–120 (cit. on pp. 23, 38).
- [Galtier et al., 2019] Mathieu N Galtier and Camille Marini. “Substra: a framework for privacy-preserving, traceable and collaborative machine learning”. In: *arXiv preprint arXiv:1910.11567* (2019) (cit. on pp. 83, 86).
- [Garcia-Dias et al., 2020] Rafael Garcia-Dias et al. “Neuroharmony: A new tool for harmonizing volumetric MRI data from unseen scanners”. In: *Neuroimage* 220 (2020) (cit. on p. 66).

- [Geyer et al., 2017] Robin C Geyer, Tassilo Klein, and Moin Nabi. “Differentially private federated learning: A client level perspective”. In: *arXiv preprint arXiv:1712.07557* (2017) (cit. on p. 17).
- [Glass, 1976] Gene V Glass. “Primary, secondary, and meta-analysis of research”. In: *Educational researcher* 5.10 (1976), pp. 3–8 (cit. on p. 9).
- [Gupta et al., 2018] Otkrist Gupta and Ramesh Raskar. “Distributed learning of deep neural network over multiple agents”. In: *Journal of Network and Computer Applications* 116 (2018), pp. 1–8 (cit. on p. 78).
- [He et al., 2020] Chaoyang He et al. “Fedml: A research library and benchmark for federated machine learning”. In: *arXiv preprint arXiv:2007.13518* (2020) (cit. on p. 83).
- [Hedges, 1992] Larry V Hedges. “Meta-analysis”. In: *Journal of Educational Statistics* 17.4 (1992), pp. 279–296 (cit. on p. 10).
- [Higgins et al., 2019] Julian PT Higgins et al. *Cochrane handbook for systematic reviews of interventions*. John Wiley & Sons, 2019 (cit. on pp. 9, 10).
- [Hsu et al., 2019] Tzu-Ming Harry Hsu, Hang Qi, and Matthew Brown. “Measuring the effects of non-identical data distribution for federated visual classification”. In: *arXiv preprint arXiv:1909.06335* (2019) (cit. on p. 30).
- [Jensen et al., 2012] Peter B Jensen, Lars J Jensen, and Søren Brunak. “Mining electronic health records: towards better research applications and clinical care”. In: *Nature Reviews Genetics* 13.6 (2012), pp. 395–405 (cit. on p. 8).
- [Johnson et al., 2007] W Evan Johnson, Cheng Li, and Ariel Rabinovic. “Adjusting batch effects in microarray expression data using empirical Bayes methods”. In: *Biostatistics* 8.1 (2007), pp. 118–127 (cit. on pp. 16, 23–26, 40, 65).
- [Jones et al., 1998a] Donald R Jones, Matthias Schonlau, and William J Welch. “Efficient Global Optimization of Expensive Black-Box Functions”. In: *Journal of Global Optimization* 13.4 (1998), pp. 455–492. URL: <https://doi.org/10.1023/A:1008306431147> (cit. on p. 54).
- [Jones et al., 1998b] Donald R Jones, Matthias Schonlau, and William J Welch. “Efficient global optimization of expensive black-box functions”. In: *Journal of Global optimization* 13.4 (1998), p. 455 (cit. on p. 55).
- [Kesteren et al., 2019] Erik-Jan van Kesteren et al. “Privacy-preserving generalized linear models using distributed block coordinate descent”. In: *arXiv preprint arXiv:1911.03183* (2019) (cit. on pp. 23, 24).

- [Kim et al., 2017] Yejin Kim et al. “Federated tensor factorization for computational phenotyping”. In: *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. 2017, pp. 887–895 (cit. on p. 78).
- [Kim et al., 2022] SeungWook Kim et al. “Harmonization of Multicenter Cortical Thickness Data by Linear Mixed Effect Model”. In: *Frontiers in Aging Neuroscience* 14 (2022) (cit. on p. 65).
- [Kingma et al., 2014] Diederik P Kingma and Jimmy Ba. “Adam: A method for stochastic optimization”. In: *arXiv preprint arXiv:1412.6980* (2014) (cit. on p. 15).
- [Konečný et al., 2016] Jakub Konečný et al. “Federated learning: Strategies for improving communication efficiency”. In: *arXiv preprint arXiv:1610.05492* (2016) (cit. on pp. 13, 15, 23, 78, 79).
- [LeCun et al., 2010] Yann LeCun and Corinna Cortes. “MNIST handwritten digit database”. In: (2010). URL: <http://yann.lecun.com/exdb/mnist/> (cit. on p. 88).
- [Lee et al., 2018] Junghye Lee et al. “Privacy-preserving patient similarity learning in a federated environment: development and analysis”. In: *JMIR medical informatics* 6.2 (2018), e20 (cit. on p. 78).
- [Leek et al., 2012] Jeffrey T Leek et al. “The sva package for removing batch effects and other unwanted variation in high-throughput experiments”. In: *Bioinformatics* 28.6 (2012), pp. 882–883 (cit. on p. 65).
- [Li et al., 2018] Tian Li et al. “Federated optimization in heterogeneous networks”. In: *arXiv preprint arXiv:1812.06127* (2018) (cit. on pp. 17, 50, 78).
- [Li et al., 2019a] Tian Li et al. “FedDane: A federated newton-type method”. In: *2019 53rd Asilomar Conference on Signals, Systems, and Computers*. IEEE. 2019, pp. 1227–1231 (cit. on pp. 50, 78).
- [Li et al., 2019b] Wenqi Li et al. “Privacy-preserving federated brain tumour segmentation”. In: *International workshop on machine learning in medical imaging*. Springer. 2019, pp. 133–141 (cit. on p. 23).
- [Li et al., 2019c] Xiang Li et al. “On the convergence of fedavg on non-iid data”. In: *arXiv preprint arXiv:1907.02189* (2019) (cit. on pp. 15, 27).
- [Li et al., 2020a] Tian Li et al. “Federated learning: Challenges, methods, and future directions”. In: *IEEE Signal Processing Magazine* 37.3 (2020), pp. 50–60 (cit. on pp. 15, 16, 40, 66, 78).
- [Li et al., 2020b] Xiaoxiao Li et al. “Multi-site fMRI analysis using privacy-preserving federated learning and domain adaptation: ABIDE results”. In: *Medical Image Analysis* 65 (2020), p. 101765 (cit. on p. 23).
- [Lim et al., 2020] Wei Yang Bryan Lim et al. “Federated learning in mobile edge networks: A comprehensive survey”. In: *IEEE Communications Surveys & Tutorials* 22.3 (2020), pp. 2031–2063 (cit. on p. 16).

- [Liu et al., 2021] Yang Liu et al. “Fate: An industrial grade platform for collaborative learning with data protection”. In: *The Journal of Machine Learning Research* 22.1 (2021), pp. 10320–10325 (cit. on pp. 83, 86).
- [Lorenzi et al., 2018] Marco Lorenzi et al. “Susceptibility of brain atrophy to TRIB3 in Alzheimer’s disease, evidence from functional prioritization in imaging genetics”. In: *Proceedings of the National Academy of Sciences* 115.12 (2018), pp. 3162–3167 (cit. on p. 29).
- [Ludwig et al., 2020] Heiko Ludwig et al. “Ibm federated learning: an enterprise framework white paper v0. 1”. In: *arXiv preprint arXiv:2007.10987* (2020) (cit. on p. 83).
- [Lydia et al., 2019] Agnes Lydia and Sagayaraj Francis. “Adagrad—an optimizer for stochastic gradient descent”. In: *Int. J. Inf. Comput. Sci* 6.5 (2019), pp. 566–568 (cit. on p. 15).
- [Ma et al., 2019] Yanjun Ma et al. “PaddlePaddle: An open-source deep learning platform from industrial practice”. In: *Frontiers of Data and Computing* 1.1 (2019), pp. 105–115 (cit. on pp. 83, 86).
- [MacKay et al., 2003] David JC MacKay and David JC Mac Kay. *Information theory, inference and learning algorithms*. Cambridge university press, 2003 (cit. on p. 52).
- [Mahlool et al., 2022] Dhurgham Hassan Mahlool and Mohamed Hamzah Abed. “Distributed brain tumor diagnosis using a federated learning environment”. In: *Bulletin of Electrical Engineering and Informatics* 11.6 (2022). Cited by: 0; All Open Access, Gold Open Access, Green Open Access, pp. 3313–3321. URL: <https://www.scopus.com/inward/record.uri?eid=2-s2.0-85139133360&doi=10.11591%2fee1.v11i6.4131&partnerID=40&md5=1ad3123e9db42a98e13e51bd26fdcbc1> (cit. on p. 23).
- [Marcus et al., 2010] Daniel S Marcus et al. “Open access series of imaging studies: longitudinal MRI data in nondemented and demented older adults”. In: *Journal of cognitive neuroscience* 22.12 (2010), pp. 2677–2684 (cit. on p. 32).
- [Marek et al., 2018] Kenneth Marek et al. “The Parkinson’s progression markers initiative (PPMI)—establishing a PD biomarker cohort”. In: *Annals of clinical and translational neurology* 5.12 (2018), pp. 1460–1477 (cit. on p. 32).
- [Mathew et al., 1999] Thomas Mathew and Kenneth Nordstrom. “On the equivalence of meta-analysis using literature and using individual patient data”. In: *Biometrics* 55.4 (1999), pp. 1221–1223 (cit. on p. 11).
- [McMahan et al., 2017] Brendan McMahan et al. “Communication-efficient learning of deep networks from decentralized data”. In: *Artificial intelligence and statistics*. PMLR. 2017, pp. 1273–1282 (cit. on pp. 13, 14, 23, 27, 50, 57, 79).

- [Merkel, 2014] Dirk Merkel. “Docker: lightweight linux containers for consistent development and deployment”. In: *Linux journal* 2014.239 (2014), p. 2 (cit. on p. 82).
- [Miotto et al., 2018] Riccardo Miotto et al. “Deep learning for healthcare: review, opportunities and challenges”. In: *Briefings in bioinformatics* 19.6 (2018), pp. 1236–1246 (cit. on p. 7).
- [Mockus, 1998] Jonas Mockus. “The application of Bayesian methods for seeking the extremum”. In: *Towards global optimization* 2 (1998), p. 117 (cit. on p. 54).
- [Murdoch et al., 2013] Travis B Murdoch and Allan S Detsky. “The inevitable application of big data to health care”. In: *Jama* 309.13 (2013), pp. 1351–1352 (cit. on p. 7).
- [Orlhac et al., 2022] Fanny Orlhac et al. “A guide to ComBat harmonization of imaging biomarkers in multicenter studies”. In: *Journal of Nuclear Medicine* 63.2 (2022), pp. 172–179 (cit. on pp. 16, 65).
- [Pardalos et al., 2021] Panos M Pardalos, Varvara Rasskazova, Michael N Vrahatis, et al. *Black Box Optimization, Machine Learning, and No-Free Lunch Theorems*. Springer, 2021 (cit. on p. 58).
- [Paszke et al., 2019] Adam Paszke et al. “Pytorch: An imperative style, high-performance deep learning library”. In: *Advances in neural information processing systems*. 2019, pp. 8026–8037 (cit. on p. 88).
- [Perrone et al., 2019] Valerio Perrone et al. *Learning search spaces for Bayesian optimization: Another view of hyperparameter transfer learning*. Tech. rep. 2019 (cit. on pp. 57, 59).
- [Plis et al., 2016] Sergey M Plis et al. “COINSTAC: a privacy enabled model and prototype for leveraging and processing decentralized brain imaging data”. In: *Frontiers in neuroscience* 10 (2016), p. 365 (cit. on p. 79).
- [Pomponio et al., 2020] Raymond Pomponio et al. “Harmonization of large MRI datasets for the analysis of brain imaging patterns throughout the lifespan”. In: *NeuroImage* 208 (2020), p. 116450 (cit. on pp. 16, 23–26, 29, 38, 43, 44).
- [Real et al., 2019] Esteban Real et al. “Regularized evolution for image classifier architecture search”. In: *Proceedings of the aaai conference on artificial intelligence*. Vol. 33. 01. 2019, pp. 4780–4789 (cit. on p. 51).
- [Reddi et al., 2020] Sashank Reddi et al. “Adaptive federated optimization”. In: *arXiv preprint arXiv:2003.00295* (2020) (cit. on p. 14).
- [Reina et al., 2021] G Anthony Reina et al. “OpenFL: An open-source framework for Federated Learning”. In: *arXiv preprint arXiv:2105.06413* (2021) (cit. on pp. 83, 87).
- [Reynolds et al., 2022] Maxwell Reynolds et al. “Combat harmonization: Empirical bayes versus fully bayes approaches”. In: *bioRxiv* (2022) (cit. on pp. 30, 33).

- [Rieke et al., 2020] Nicola Rieke et al. “The future of digital health with federated learning”. In: *NPJ digital medicine* 3.1 (2020), pp. 1–7 (cit. on p. 50).
- [Rosenthal et al., 2016] Liana S Rosenthal et al. “The NINDS Parkinson’s disease biomarkers program”. In: *Movement Disorders* 31.6 (2016), pp. 915–923 (cit. on p. 32).
- [Roth et al., 2020] Holger R Roth et al. “Federated learning for breast density classification: A real-world implementation”. In: *Domain Adaptation and Representation Transfer, and Distributed and Collaborative Learning: Second MICCAI Workshop, DART 2020, and First MICCAI Workshop, DCL 2020, Held in Conjunction with MICCAI 2020, Lima, Peru, October 4–8, 2020, Proceedings 2*. Springer. 2020, pp. 181–191 (cit. on p. 87).
- [Roth et al., 2022] Holger R Roth et al. “NVIDIA FLARE: Federated Learning from Simulation to Real-World”. In: *arXiv preprint arXiv:2210.13291* (2022) (cit. on p. 83).
- [Rumsfeld et al., 2016] John S Rumsfeld, Karen E Joynt, and Thomas M Maddox. “Big data analytics to improve cardiovascular care: promise and challenges”. In: *Nature Reviews Cardiology* 13.6 (2016), pp. 350–359 (cit. on p. 8).
- [Ryffel et al., 2018] Theo Ryffel et al. “A generic framework for privacy preserving deep learning”. In: *arXiv preprint arXiv:1811.04017* (2018) (cit. on p. 80).
- [Sarma et al., 2021] Karthik V Sarma et al. “Federated learning improves site performance in multicenter deep learning without data sharing”. In: *Journal of the American Medical Informatics Association* 28.6 (2021), pp. 1259–1264 (cit. on p. 87).
- [Satizabal et al., 2019] Claudia L Satizabal et al. “Genetic architecture of subcortical brain structures in 38,851 individuals”. In: *Nature genetics* 51.11 (2019), pp. 1624–1636 (cit. on p. 11).
- [Sattler et al., 2019] Felix Sattler et al. “Robust and communication-efficient federated learning from non-iid data”. In: *IEEE transactions on neural networks and learning systems* 31.9 (2019), pp. 3400–3413 (cit. on p. 17).
- [Schmaal et al., 2016] Lianne Schmaal et al. “Subcortical brain alterations in major depressive disorder: findings from the ENIGMA Major Depressive Disorder working group”. In: *Molecular psychiatry* 21.6 (2016), pp. 806–812 (cit. on p. 11).
- [Schneeweiss, 2014] Sebastian Schneeweiss. “Learning from big health care data”. In: *N Engl J Med* 370.23 (2014), pp. 2161–2163 (cit. on p. 7).

- [Sheller et al., 2019] Micah J Sheller et al. “Multi-institutional deep learning modeling without sharing patient data: A feasibility study on brain tumor segmentation”. In: *Brainlesion: Glioma, Multiple Sclerosis, Stroke and Traumatic Brain Injuries: 4th International Workshop, BrainLes 2018, Held in Conjunction with MICCAI 2018, Granada, Spain, September 16, 2018, Revised Selected Papers, Part I* 4. Springer. 2019, pp. 92–104 (cit. on p. 15).
- [Shen et al., 2016] Shiqi Shen, Shruti Tople, and Prateek Saxena. “Auror: Defending against poisoning attacks in collaborative deep learning systems”. In: *Proceedings of the 32nd Annual Conference on Computer Security Applications*. 2016, pp. 508–519 (cit. on p. 83).
- [Silva et al., 2019] Santiago Silva et al. “Federated learning in distributed medical databases: Meta-analysis of large-scale subcortical brain data”. In: *2019 IEEE 16th international symposium on biomedical imaging (ISBI 2019)*. IEEE. 2019, pp. 270–274 (cit. on pp. 27, 29, 80, 90).
- [Sled et al., 1998] John G Sled, Alex P Zijdenbos, and Alan C Evans. “A nonparametric method for automatic correction of intensity nonuniformity in MRI data”. In: *IEEE transactions on medical imaging* 17.1 (1998), pp. 87–97 (cit. on p. 32).
- [Smith et al., 2017] Virginia Smith et al. “Federated multi-task learning”. In: *Advances in neural information processing systems* 30 (2017) (cit. on p. 16).
- [Solove et al., 2020] Daniel J Solove and Paul M Schwartz. *Information privacy law*. Aspen Publishing, 2020 (cit. on pp. 8, 9, 11).
- [Sperling et al., 2014] Reisa A Sperling et al. “The A4 study: stopping AD before symptoms begin?” In: *Science translational medicine* 6.228 (2014), 228fs13–228fs13 (cit. on p. 32).
- [Sun et al., 2013] Jimeng Sun and Chandan K Reddy. “Big data analytics for healthcare”. In: *Proceedings of the 19th ACM SIGKDD international conference on Knowledge discovery and data mining*. 2013, pp. 1525–1525 (cit. on p. 8).
- [Thompson et al., 2014] Paul M Thompson et al. “The ENIGMA Consortium: large-scale collaborative analyses of neuroimaging and genetic data”. In: *Brain imaging and behavior* 8.2 (2014), pp. 153–182 (cit. on p. 93).
- [Torous et al., 2018] John Torous et al. “Clinical review of user engagement with mental health smartphone apps: evidence, theory and improvements”. In: *BMJ Ment Health* 21.3 (2018), pp. 116–119 (cit. on p. 7).
- [Undavia et al., 2020] Jaimin Navinchandra Undavia and Atul Manubhai Patel. “Big Data Analytics in Healthcare: Applications and Challenges”. In: *International Journal of Big Data and Analytics in Healthcare (IJBDAH)* 5.1 (2020), pp. 19–27 (cit. on p. 7).

- [Voigt et al., 2017] Paul Voigt and Axel Von dem Bussche. “The eu general data protection regulation (gdpr)”. In: *A Practical Guide, 1st Ed., Cham: Springer International Publishing* (2017) (cit. on pp. 8, 11, 78).
- [Wachinger et al., 2021] Christian Wachinger et al. “Detect and correct bias in multi-site neuroimaging datasets”. In: *Medical Image Analysis* 67 (2021), p. 101879 (cit. on pp. 16, 22).
- [Webb-Vargas et al., 2017] Yenny Webb-Vargas et al. “Big data and neuroimaging”. In: *Statistics in biosciences* 9 (2017), pp. 543–558 (cit. on pp. 8, 9).
- [Weiner et al., 2013] Michael W Weiner et al. “The Alzheimer’s Disease Neuroimaging Initiative: a review of papers published since its inception”. In: *Alzheimer’s & Dementia* 9.5 (2013), e111–e194 (cit. on p. 32).
- [Wosik et al., 2020] Jedrek Wosik et al. “Telehealth transformation: COVID-19 and the rise of virtual care”. In: *Journal of the American Medical Informatics Association* 27.6 (2020), pp. 957–962 (cit. on p. 7).
- [Yang et al., 2019] Qiang Yang et al. “Federated machine learning: Concept and applications”. In: *ACM Transactions on Intelligent Systems and Technology (TIST)* 10.2 (2019), pp. 1–19 (cit. on p. 78).
- [Yuhong et al., 2019] Wen Yuhong et al. *Federated Learning powered by NVIDIA Clara*. 2019. URL: <https://devblogs.nvidia.com/federated-learning-clara/> (visited on June 22, 2020) (cit. on p. 79).
- [Zaheer et al., 2018] Manzil Zaheer et al. “Adaptive methods for nonconvex optimization”. In: *Advances in neural information processing systems* 31 (2018) (cit. on p. 15).
- [Zhao et al., 2018] Yue Zhao et al. “Federated learning with non-iid data”. In: *arXiv preprint arXiv:1806.00582* (2018) (cit. on p. 16).
- [Ziller et al., 2021] Alexander Ziller et al. “Pysyft: A library for easy federated learning”. In: *Federated Learning Systems: Towards Next-Generation AI* (2021), pp. 111–139 (cit. on pp. 83, 84).
- [Zugman et al., 2022] André Zugman et al. “Mega-analysis methods in ENIGMA: The experience of the generalized anxiety disorder working group”. In: *Human Brain Mapping* 43.1 (2022), pp. 255–277 (cit. on p. 11).
- [Zuo et al., 2021] Lianrui Zuo et al. “Unsupervised MR harmonization by learning disentangled representations using information bottleneck theory”. In: *NeuroImage* 243 (2021), p. 118569 (cit. on p. 66).