



HAL
open science

Fed-MIWAE: Federated Imputation of Incomplete Data via Deep Generative Models

Irene Balelli, Aude Sportisse, Francesco Cremonesi, Pierre-Alexandre Mattei,
Marco Lorenzi

► **To cite this version:**

Irene Balelli, Aude Sportisse, Francesco Cremonesi, Pierre-Alexandre Mattei, Marco Lorenzi. Fed-MIWAE: Federated Imputation of Incomplete Data via Deep Generative Models. 2023. hal-04069795

HAL Id: hal-04069795

<https://hal.science/hal-04069795v1>

Preprint submitted on 14 Apr 2023

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.

Fed-MIWAE: Federated Imputation of Incomplete Data via Deep Generative Models

Irene Balelli^{1†}[0000-0002-4593-8217], Aude Sportisse^{1†}[0000-0003-1303-0762],
Francesco Cremonesi¹[0000-0003-1027-485X], Pierre-Alexandre
Mattei¹[0000-0002-1297-908X], and Marco Lorenzi¹[0000-0003-0521-2881] for the
Alzheimer’s Disease Neuroimaging Initiative*

Centre Inria d’Université Côte d’Azur

{irene.balelli,aude.sportisse,francesco.cremonesi,pierre-alexandre.mattei,marco.lorenzi}@inria.fr

Abstract. Federated learning allows for the training of machine learning models on multiple decentralized local datasets without requiring explicit data exchange. However, data pre-processing, including strategies for handling missing data, remains a major bottleneck in real-world federated learning deployment, and is typically performed locally. This approach may be biased, since the subpopulations locally observed at each center may not be representative of the overall one. To address this issue, this paper first proposes a more consistent approach to data standardization through a federated model. Additionally, we propose Fed-MIWAE, a federated version of the state-of-the-art imputation method MIWAE, a deep latent variable model for missing data imputation based on variational autoencoders. MIWAE has the great advantage of being easily trainable with classical federated aggregators. Furthermore, it is able to deal with MAR (Missing At Random) data, a more challenging missing-data mechanism than MCAR (Missing Completely At Random), where the missingness of a variable can depend on the observed ones. We evaluate our method on multi-modal medical imaging data and clinical scores from a simulated federated scenario with the ADNI dataset. We compare Fed-MIWAE with respect to classical imputation methods, either performed locally or in a centralized fashion. Fed-MIWAE allows to achieve imputation accuracy comparable with the best centralized method, even when local data distributions are highly heterogeneous. In addition, thanks to the variational nature of Fed-MIWAE, our method is designed to perform multiple imputation, allowing for the quantification of the imputation uncertainty in the federated scenario.

Keywords: Missing data · Federated learning · Federated pre-processing · Variational autoencoders · Deep Learning.

[†]The two authors contributed equally to this paper.

*Data used in preparation of this article were obtained from the Alzheimer’s Disease Neuroimaging Initiative (ADNI) database (adni.loni.usc.edu). As such, 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: http://adni.loni.usc.edu/wp-content/uploads/how_to_apply/ADNIAcknowledgement_List.pdf

1 Introduction and previous work

Digital health revolution relies heavily on the massive development and deployment of data-driven machine learning (ML) and deep learning (DL) methods. Nevertheless, in order to insure reliability, trustworthiness and generalizability of such models, their parameters need to be learned from a sufficiently large curated dataset, representative of the whole data distribution. These conditions are rarely met within a single research hospital, and data centralization is often not possible due to current privacy regulations (*e.g.* the European GDPR¹).

Federated learning (FL) is an emerging learning paradigm based on multiple decentralized datasets. FL has been identified as a promising solution for health-care applications [Rieke et al., 2020, Crowson et al., 2022], due to its improved robustness and generalization properties. Indeed, the final aggregated FL model gathers information provided by multiple hospitals through distributed optimization. Moreover, FL implicitly alleviates data governance burden and naturally improves privacy guarantees with respect to classical centralized ML methods. The growing interest in FL is demonstrated by the several FL-based research projects [Sadilek et al., 2021, Crowson et al., 2022, Rieke et al., 2020, Pati et al., 2022], and FL infrastructures such as Flare [Roth et al.], OpenFL [Foley et al.], SubstraFL², and Fed-BioMed³, which are specifically tailored for biomedical applications. Despite the potential of FL for cross-institutional medical research, data pre-processing, including data standardization and management of missing entries in the database, is generally carried out locally by each participating site. The importance of establishing reliable data pre-processing pipelines in a federated manner has already been emphasized [Rajula et al., 2019], and is of paramount importance in real-life FL deployment.

In order to prevent the deletion of incomplete rows in the dataset, with consequent critical loss of information (especially in the high-dimensional setting) and the potential induction of a biased analysis, one common approach to handle missing values involves data imputation [Little and Rubin, 2019]. Besides the naïve mean imputation, which consists in replacing missing values with the mean across the observed population, several alternatives and more elaborate imputation methods have been proposed in the literature, including based on probabilistic models [Little and Rubin, 2019], low-rank models [Mazumder et al., 2010] or random forests [Stekhoven and Bühlmann, 2012]. Some methods are specifically designed for multiple imputation, such as MICE [Van Buuren and Groothuis-Oudshoorn, 2011], to account for imputation uncertainty Austin et al. [2021]. More recently, imputation methods based on deep latent variable models have been proposed, for example by building on generative adversarial networks (GANs) [Li et al., 2019, Yoon et al., 2018] or on variational autoencoders (VAEs) [Mattei and Frellsen, 2019, Nazabal et al., 2020, Ipsen et al., 2021]. However, none of these imputations methods are specifically designed for

¹<https://gdpr-info.eu/>

²https://docs.substra.org/en/stable/substrafll_doc/substrafll_overview.html

³<https://fedbiomed.gitlabpages.inria.fr>

the FL setting. The theoretical behaviour of learning with hidden variables in a federated context has been addressed in Dieuleveut et al. [2021] through an Expectation Maximization (EM) based optimization scheme, and applied to federated missing values imputation for biodiversity monitoring. In [Balelli et al., 2021], authors proposed a Bayesian latent variable model for multiple views data assimilation and completely missing modalities imputation in a federated setting [Balelli et al., 2021], which also relies on the EM algorithm for local parameters’ optimization. Recently, two papers [Zhou et al., 2021, Yao and Zhao, 2022] have extended the imputation method based on GANs [Yoon et al., 2018] to the FL framework. Nevertheless, the theoretical results behind these methods have only been provided in the Missing Completely At Random (MCAR) setting [Yoon et al., 2018], and are not suited for the more complex and yet relevant case of Missing At Random (MAR) data, where the missingness of a variable depends on the observed ones, or the most challenging Missing Not At Random (MNAR) setting, where the missingness can even depend on the unobserved variable itself [Mohan, 2018, Austin et al., 2021] (Table 1). It is worth mentioning that it is not possible to assess whether the missing pattern is MCAR or MAR solely based on the data (*i.e.* without additional clinical knowledge), hence the need in favouring methods able to handle both cases, and possibly explicitly account for the uncertainty on the imputed variable, for instance through multiple imputation.

Table 1: Missing data classification with examples

MCAR	Accidental loss/damage of a sample, non-systematic personnel error in a sample’s acquisition/reporting.
MAR	The acquisition of a feature (<i>e.g.</i> a clinical score test) depends on a fully observed variable (<i>e.g.</i> the age): the observed systematic difference in incomplete features can be explained by differences in observed variables.
MNAR	The reason for a missing sample can depend both on observed and unobserved variables, including the missing sample itself (<i>e.g.</i> self-censored data)

In this work, we first propose a pipeline for data pre-processing through multi-site data, including federated standardization as well as imputation of missing entries, performed through Fed-MIWAE, the federated version of MIWAE [Mattei and Frellsen, 2019], a deep latent variable model for missing data imputation based on variational inference. The VAEs-based structure of MIWAE makes it easily trainable with classical federated aggregators, such as FedAvg [McMahan et al., 2017], FedProx [Li et al., 2018] or Scaffold [Karimireddy et al., 2020]. By relying on the framework of MIWAE, Fed-MIWAE is naturally designed to handle both MCAR and MAR data, while allowing for multiple imputation. In Sec. 2 we describe our proposed approach and its implementation, and in Sec. 3 we applied it to multi-modal medical imaging data and clinical scores from the ADNI database, distributed to multiple sites to simulate a federated setting. Fed-MIWAE shows higher generalization and improved performance compared to all tested local imputation methods, and comparable with the centralized sce-

nario, even in presence of simulated MAR data and high heterogeneity across local datasets in terms of patients distribution.

2 An approach to federated data pre-processing

2.1 Missing data importance-weighted autoencoder (MIWAE)

In this work we extend the MIWAE method [Mattei and Frellsen, 2019] in order to be able to train it in a federated manner. MIWAE is defined by a deep generative model from a latent Gaussian variable z to the original data x , where $t \sim p_\theta(x|z)$, and the conditional distribution p_θ is parametrized using a neural network (NN) with weights θ . Due to the presence of missing values in the data, the observed likelihood we would like to maximise is intractable [Little and Rubin, 2019]. Similarly to variational autoencoders, the idea of MIWAE is to propose a conditional distribution of the latent variables given the observed data, x^{observed} , $q_\gamma(z|x^{\text{observed}})$, parameterized as well by a NN with weights γ . A stochastic lower bound of the observed likelihood can then be computed and optimized using SGD-like approaches. Finally, the imputation step is done by computing the conditional expectation of the missing variables given the observed ones, $\mathbb{E}[x^{\text{missing}}|x^{\text{observed}}]$, estimated via self-normalised importance sampling.

2.2 Proposed method

We propose a federated pre-processing approach, described in Algorithm 1. Firstly, we perform federated standardization in presence of missing values through a single communication round (Stage 1). This step only requires the communication to the central server of all local feature-wise means and standard deviations to analytically evaluate the global mean and std. Each center will be able to use this information to perform its data standardization with respect to the whole federated dataset. Secondly, we propose to train Fed-MIWAE using the same FL platform (Stage 2 (i) of Algorithm 1), over the decentralised standardised datasets, across R iterative communication rounds. In this stage, the local parameters' optimization and aggregation can be performed using classical FL optimization and aggregation schemes, such as Federated Averaging (FedAvg) [McMahan et al., 2017], FedProx [Li et al., 2020] or Scaffold [Karimireddy et al., 2020], which are designed to be more robust to data heterogeneity. Finally, missing values can be imputed locally by each center using the final global model (Stage 2 (ii) of Algorithm 1).

3 Application to the Alzheimer's Disease Neuroimaging Initiative dataset (ADNI)

3.1 Data

We consider 311 participants extracted from the ADNI dataset, composed by 104 cognitively normal individuals (NC) and 207 patients diagnosed with Alzheimer

Algorithm 1 Federated data pre-processing

Input: Rounds R , Epochs E , MIWAE hyperparameters**Stage 1: Federated standardization**Each center c computes its local mean and std $(\boldsymbol{\mu}^{(c)}, \boldsymbol{\sigma}^{(c)})$.Server aggregates them and sends the global mean and std $(\boldsymbol{\mu}, \boldsymbol{\sigma})$.Each center c standardizes its local dataset with $\boldsymbol{\mu}$ and $\boldsymbol{\sigma}$.**Stage 2: Fed-MIWAE****(i) Train step****for** $r = 1$ **to** R **do** Each center c initializes its local model with $(\theta^{r-1}, \gamma^{r-1})$. Each center performs E epochs of local MIWAE bound optimization, and sends the updated local parameters $((\theta^{(c)})^{t+1}, (\gamma^{(c)})^{t+1})$ to the Server. Server aggregates $((\theta^{(c)})^{t+1}, (\gamma^{(c)})^{t+1})_c$ and sends the global parameter (θ^r, γ^r) back to the centers.**end for****(ii) Missing values imputation**Each center c dispose of the final global model (θ^R, γ^R) for missing data imputation.

disease (AD), coming from 55 centers. All participants are associated with multiple data views: cognitive scores including MMSE, CDR-SB, ADAS-Cog-11 and RAVLT, Magnetic resonance imaging (MRI), Fluorodeoxyglucose-PET and AV45-Amyloid PET images. MRI morphometrical biomarkers were obtained as regional volumes using the cross-sectional pipeline of FreeSurfer v6.0 based on the Desikan-Killiany parcellation [Fischl, 2012]. Measurements from AV45-PET and FDG-PET were estimated by co-registering each modality to their respective MRI, normalizing by the cerebellum uptake and by computing regional amyloid load and glucose hypometabolism using PetSurfer pipeline [Greve et al., 2014] based on the same parcellation. Features were corrected beforehand with respect to intra-cranial volume, sex and age using a multivariate linear model. Overall, the features dimension for a fully observed patient is 130. Further details are provided in Supp. Material, Table S1.

3.2 Scenarios

We consider two federated scenarios:

Natural split scenario: we take into account the information concerning the site to perform the split. Due to the high number of participating sites, we artificially create 4 datasets by gathering the contribution of about 14 ADNI sites per simulated center (see details in Supp. Material, Table S2).

Not-IID scenario: we generated 3 centers, one containing only AD patients (Client 1, $N_1 = 171$), one containing only NC patients (Client 2, $N_2 = 77$), and the third one, kept as external testing dataset, containing samples from both groups ($N_3 = 63$).

We further simulate 30% of missing data in each resulting local dataset, considering either the MCAR or MAR settings. For the natural split scenario, we

perform 5 times a 4-fold cross-validation: each time, three datasets are used as the training centers, while the fourth one is used in testing phase only. For the not-IID scenario, the training is always performed (5 times) on Client 1 and Client 2, with testing on Client 3. Numerical experiments with Fed-MIWAE were executed using version v4.2 of the Fed-BioMed software⁴. Experiment results were logged using Fed-BioMed’s logging facility and manually collected for analysis. Due to space limitations, we will only show results on the more challenging MAR setting in Sec. 3.4. Similar conclusions have been obtained in the MCAR setting, and are provided in the Supp. Materials (Table S5, Figure S1).

3.3 Benchmark

We compare our proposed method trained using FedProx [Li et al., 2020] (simply denoted **FedProx** in Sec. 3.4, or **FedProx_loc**, if the data standardization is done locally and only Stage 2 of Algorithm 1 is performed), with the local training of MIWAE in each site (**Local_cli** for the site i) as well as the centralized training of MIWAE (**Centralized**). To ensure a fair comparison, the chosen hyperparameters of MIWAE are the same, regardless of how the model is trained (Supp. Material, Table S3). In addition, we assess the performance of classical imputation methods, by training them either locally or using the centralized dataset: naïve mean imputation (**mean**), nonparametric method with random forests (**rf**) Stekhoven and Bühlmann [2012] and iterative chained equations method (**ice**) [Van Buuren and Groothuis-Oudshoorn, 2011], from the impute module of the scikit-learn library Pedregosa et al. [2011]. To evaluate the imputation accuracy of all considered methods, we compute the normalized Mean Squared Error (MSE).

3.4 Results

In Figure 1, we assess the performance of the imputation methods through MSE in both considered scenarios, with respect to the task of imputing missing data both on the external testing dataset (upper row) and on datasets which have been seen during training (bottom row). In all cases, our method optimized with FedProx outperforms all local MIWAE models (**Local_cli**). Moreover, one can see that performing data standardisation in a federated manner allows to reach improved results compared to the case where the standardization step has been performed locally (*cf* **FedProx** compared to **FedProx_loc**), especially in the more challenging non-iid scenario. Interestingly, the MIWAE method in the centralized setting displays a lower performance compared to Fed-MIWAE. This can be attributed to the potential issue of overfitting, and can be prevented for instance by deploying drop-out and early stopping strategies, which has not been done here for the sake of a fair comparison throughout the methods.

Among the considered benchmark approaches, one can see that the classical Random Forests imputation method (**rf**) achieves the highest performance,

⁴<https://gitlab.inria.fr/fedbiomed/fedbiomed/-/tree/v4.2.4>

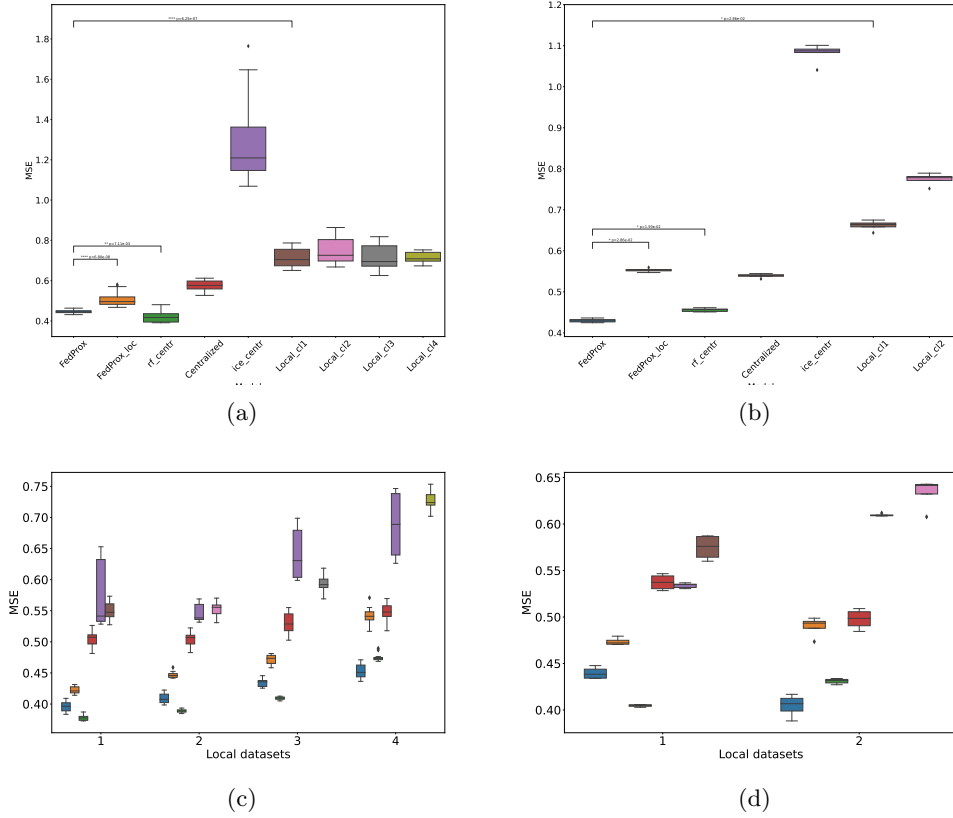


Fig. 1: ADNI data, MAR setting. Comparison of federated and local models measured on imputing missing entries of a previously unseen testing dataset (upper row) and of local datasets used for training (bottom row, where colors denotes the employed model, following the same color code as in the upper row figures). Left column: Natural split scenario. Right column: Not-IID scenario. More details can be found in Table 2 and Supp. Table S4

when tested on both an external dataset and the completion of missing entries from the training data. The naive mean-values imputation method `mean` is not included in Figure 1 due to its very poor performance (see for instance Table 2), while the method `ice`, is significantly less accurate than `rf` and our proposed approach, despite being trained in a standard centralized manner.

Fed-MIWAE shows a very good performance in all considered settings, compared to the centralized approaches tested. In particular, the imputation quality of our federated method is remarkable when considering the not-IID scenario and the ability of the trained model to infer missing entries in an external dataset (Figure 1 (d) and Table 2), outperforming the best centralized model. In addi-

tion, one can appreciate the consistent performance exhibited by Fed-MIWAE in the not-IID setting, where it achieves a MSE similar to the one obtained in the natural split scenario, where patients’ distribution across centers is more balanced. In contrast, the performance of all local models (`Local_cli`) in imputing missing values of the external testing dataset drops significantly in this extremely non-heterogeneous scenario. Finally, our approach is even able to outperform the best centralized model when applied to the smaller sized local datasets (*i.e.* Client 4 in the natural split scenario and Client 2 in the non-iid one), underscoring the importance of performing the imputation task by exploiting the federated framework.

Table 2: ADNI data, MAR setting, not-IID split. Comparison of all considered methods on the testing dataset (mean (std)).

Training	Federated		Centralized		
Model	FedProx	FedProx_loc	mean	ice	random forests
MSE ↓	0.4299 (0.005)	0.5529 (0.005)	1.0933	1.081 (0.0233)	0.4559 (0.0042)

Finally, we should stress that an additional major advantage of our method consists in its ability to perform multiple imputation, by drawing i.i.d. samples from the estimated posterior probability $p_{\text{Fed-MIWAE}}(x^{\text{missing}}|x^{\text{observed}})$. This allows us to quantify the uncertainty of the imputation. In Figure 2 (a) we show an example of multiple imputation on a randomly selected subject in the testing dataset obtained by drawing 30 samples, while in Figure 2 (b) one can see that a significantly lower variance is associated to $p_{\text{Fed-MIWAE}}(x^{\text{missing}}|x^{\text{observed}})$ compared to $p(x^{\text{missing}})$, leading us to the conclusion that x^{observed} do provide relevant information for the imputation of x^{missing} .

4 Conclusion

In this work, we propose an efficient pipeline to handle MCAR and MAR values in the FL framework, which provides highly generalizable results and comparable performance with respect to centralized methods. We show that our method is accurate even when applied in highly heterogeneous settings, which should be carefully taken into account in real-world FL contexts. Several improvement of the proposed method can be envisaged in the future, spanning from addressing the complex problem of handling MNAR data, by incorporating the not-MIWAE method [Ipsen et al., 2021] in our pipeline, to improve privacy requirements for instance by using secure aggregation or differential privacy techniques. In addition, data pre-processing pipelines in the federated framework should be further explored since they are key for the successful training of FL in real clinical studies.

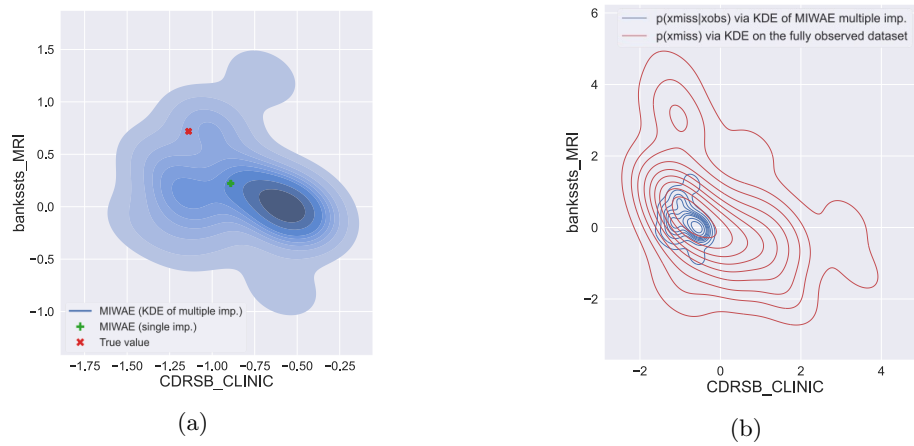


Fig. 2: ADNI data, MAR setting. (a) Example of multiple imputation on the test data using Algorithm 1, and (b) corresponding comparison between $p_{\text{Fed-MIWAE}}(x^{\text{missing}}|x^{\text{observed}})$ with $p(x^{\text{missing}})$.

Bibliography

- Peter C Austin, Ian R White, Douglas S Lee, and Stef van Buuren. Missing data in clinical research: a tutorial on multiple imputation. Canadian Journal of Cardiology, 37(9):1322–1331, 2021.
- Irene Balelli, Santiago Silva, Marco Lorenzi, Alzheimer’s Disease Neuroimaging Initiative, et al. A probabilistic framework for modeling the variability across federated datasets. In International Conference on Information Processing in Medical Imaging, pages 701–714. Springer, 2021.
- Matthew G Crowson, Dana Moukheiber, Aldo Robles Arévalo, Barbara D Lam, Sreekar Mantena, Aakanksha Rana, Deborah Goss, David W Bates, and Leo Anthony Celi. A systematic review of federated learning applications for biomedical data. PLOS Digital Health, 1(5):e0000033, 2022.
- Aymeric Dieuleveut, Gersende Fort, Eric Moulines, and Geneviève Robin. Federated-EM with heterogeneity mitigation and variance reduction. In Advances in Neural Information Processing Systems, volume 34, pages 29553–29566. Curran Associates, Inc., 2021.
- Bruce Fischl. Freesurfer. Neuroimage, 62(2):774–781, 2012.
- Patrick Foley, Micah J Sheller, Brandon Edwards, Sarthak Pati, Walter Riviera, Mansi Sharma, Prakash Narayana Moorthy, Shi-han Wang, Jason Martin, Parsa Mirhaji, Prashant Shah, and Spyridon Bakas. Openfl: the open federated learning library. Physics in Medicine & Biology.
- Douglas N Greve, Claus Svarer, Patrick M Fisher, Ling Feng, Adam E Hansen, William Baare, Bruce Rosen, Bruce Fischl, and Gitte M Knudsen. Cortical surface-based analysis reduces bias and variance in kinetic modeling of brain pet data. Neuroimage, 92:225–236, 2014.
- Niels Bruun Ipsen, Pierre-Alexandre Mattei, and Jes Frelsen. not-MIWAE: Deep generative modelling with missing not at random data. In International Conference on Learning Representations, ICLR, 2021.
- Sai Praneeth Karimireddy, Satyen Kale, Mehryar Mohri, Sashank Reddi, Sebastian Stich, and Ananda Theertha Suresh. Scaffold: Stochastic controlled averaging for federated learning. In International Conference on Machine Learning, pages 5132–5143. PMLR, 2020.
- Steven Cheng-Xian Li, Bo Jiang, and Benjamin Marlin. Learning from incomplete data with generative adversarial networks. 2019.
- Tian Li, Anit Kumar Sahu, Manzil Zaheer, Maziar Sanjabi, Ameet Talwalkar, and Virginia Smith. Federated optimization in heterogeneous networks. arXiv preprint arXiv:1812.06127, 2018.
- Tian Li, Anit Kumar Sahu, Ameet Talwalkar, and Virginia Smith. Federated learning: Challenges, methods, and future directions. IEEE Signal Processing Magazine, 37(3):50–60, 2020.
- Roderick JA Little and Donald B Rubin. Statistical analysis with missing data, volume 793. John Wiley & Sons, 2019.

- Pierre-Alexandre Mattei and Jes Frellsen. MIWAE: Deep generative modelling and imputation of incomplete data sets. In International conference on machine learning, pages 4413–4423. PMLR, 2019.
- Rahul Mazumder, Trevor Hastie, and Robert Tibshirani. Spectral regularization algorithms for learning large incomplete matrices. The Journal of Machine Learning Research, 11:2287–2322, 2010.
- Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Aguera y Arcas. Communication-efficient learning of deep networks from decentralized data. In Artificial intelligence and statistics, pages 1273–1282. PMLR, 2017.
- Karthika Mohan. On handling self-masking and other hard missing data problems. In AAAI symposium, 2018.
- Alfredo Nazabal, Pablo M Olmos, Zoubin Ghahramani, and Isabel Valera. Handling incomplete heterogeneous data using vaes. Pattern Recognition, 107:107501, 2020.
- Sarthak Pati, Ujjwal Baid, Brandon Edwards, Micah Sheller, Shih-Han Wang, G Anthony Reina, Patrick Foley, Alexey Gruzdev, Deepthi Karkada, Christos Davatzikos, et al. Federated learning enables big data for rare cancer boundary detection. Nature communications, 13(1):7346, 2022.
- F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. Journal of Machine Learning Research, 12:2825–2830, 2011.
- Hema Sekhar Reddy Rajula, Veronika Odintsova, Mirko Manchia, and Vassilios Fanos. Overview of federated facility to harmonize, analyze and management of missing data in cohorts. Applied Sciences, 9(19):4103, 2019.
- Nicola Rieke, Jonny Hancox, Wenqi Li, Fausto Milletari, Holger R Roth, Shadi Albarqouni, Spyridon Bakas, Mathieu N Galtier, Bennett A Landman, Klaus Maier-Hein, et al. The future of digital health with fl. NPJ digital medicine, 3(1):1–7, 2020.
- Holger R Roth, Yan Cheng, Yuhong Wen, Isaac Yang, Ziyue Xu, YuanTing Hsieh, Kristopher Kersten, Ahmed Harouni, Can Zhao, Kevin Lu, et al. Flare: Federated learning from simulation to real-world. In Workshop on Federated Learning: Recent Advances and New Challenges (in Conjunction with NeurIPS 2022).
- Adam Sadilek, Luyang Liu, Dung Nguyen, Methun Kamruzzaman, Stylianos Serghiou, Benjamin Rader, Alex Ingerman, Stefan Mellem, Peter Kairouz, Elaine O Nsoesie, et al. Privacy-first health research with federated learning. NPJ digital medicine, 4(1):1–8, 2021.
- Daniel J Stekhoven and Peter Bühlmann. Missforest—non-parametric missing value imputation for mixed-type data. Bioinformatics, 28(1):112–118, 2012.
- Stef Van Buuren and Karin Groothuis-Oudshoorn. mice: Multivariate imputation by chained equations in r. Journal of statistical software, 45:1–67, 2011.
- Zoujing Yao and Chunhui Zhao. Fedtmi: Knowledge aided federated transfer learning for industrial missing data imputation. Journal of Process Control, 117:206–215, 2022.

Jinsung Yoon, James Jordon, and Mihaela Schaar. Gain: Missing data imputation using generative adversarial nets. In International conference on machine learning, pages 5689–5698. PMLR, 2018.

Xu Zhou, Xiaofeng Liu, Gongjin Lan, and Jian Wu. Federated conditional generative adversarial nets imputation method for air quality missing data. Knowledge-Based Systems, 228:107261, 2021.

Supplementary Material

Table S1: Demographics of the clinical sample from ADNI.

Group	Sex	Count	Age	Range
AD	Female	94	71.58 (7.59)	55.10 - 90.30
	Male	113	74.37 (7.19)	55.90 - 89.30
NC	Female	58	73.76 (4.61)	65.10 - 84.70
	Male	46	75.39 (6.58)	59.90 - 85.60

Table S2: Clinical sample’s sizes per simulated center, natural split.

Center	1	2	3	4
Count	AD 57 NC 35	AD 66 NC 38	AD 41 NC 21	AD 43 NC 10
Total	92	104	62	53

Table S3: Fed-MIWAE Hyperparameters used. C = Number of clients.

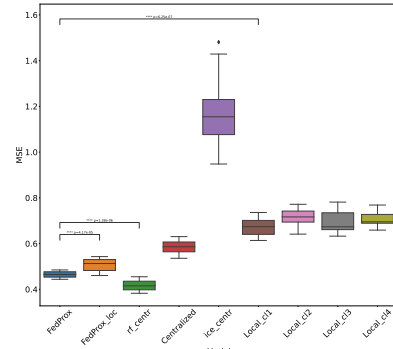
Training	Rounds	Epochs	Hidden units	Latent dim.	N samples (train)	N samples (test)
local_cli	1	150×10	256	20	50	10000
centralized	1	$150 \times 10 \times C$				
federated	150	10				

Table S4: ADNI data, MAR setting, Natural split. Testing dataset.

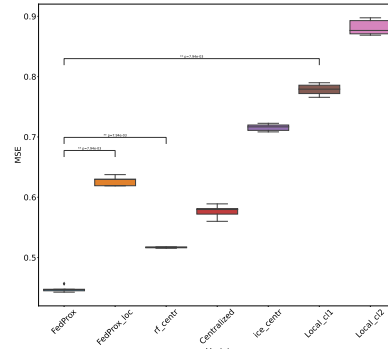
Training	Federated		Centralized		
Model	FedProx	FedProx_loc	mean	ice	random forests
MSE ↓	0.4458 (0.0083)	0.5086 (0.0383)	1.117 (0.0555)	1.294 (0.2165)	0.424 (0.0338)
mean (std)					

Table S5: ADNI data, MCAR setting, not-IID split. Testing dataset.

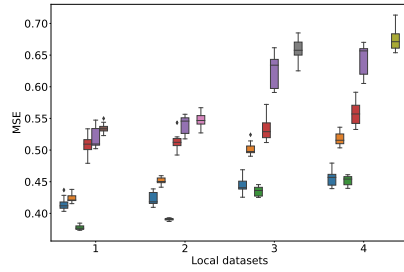
Training	Federated		Centralized		
Model	FedProx	FedProx_loc	mean	ice	random forests
MSE	0.4477 (0.0053)	0.627 (0.008)	1.1717	0.7157 (0.006)	0.5167 (0.0009)
mean (std)					



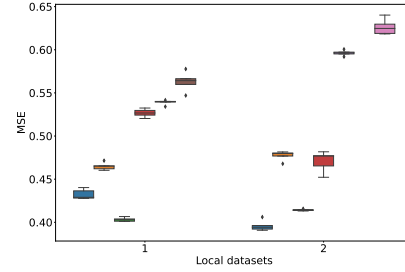
(a)



(b)



(c)



(d)

Fig. S1: ADNI data, MCAR setting. Comparison of federated and local models measured on imputing missing entries of a previously unseen testing dataset (upper row) and of local datasets used for training (bottom row, where colors denotes the employed model, following the same color code as in the upper row figures). Left column: Natural split scenario. Right column: Not-IID scenario.