

FedHarmony: Unlearning Scanner Bias with Distributed Data

Nicola K. Dinsdale¹, Mark Jenkinson^{2,3,4}, and Ana I.L. Namburete¹

¹ Oxford Machine Learning in NeuroImaging (OMNI) Lab, University of Oxford, UK

² Wellcome centre for Integrative Neuroimaging, FMRIB, University of Oxford, Oxford, UK

³ Australian Institute for Machine Learning (AIML), Department of Computer Science, University of Adelaide, Adelaide, Australia

⁴ South Australian Health and Medical Research Institute (SAHMRI), North Terrace, Adelaide, Australia
`nicola.dinsdale@cs.ox.ac.uk`

Abstract. The ability to combine data across scanners and studies is vital for neuroimaging, to increase both statistical power and the representation of biological variability. However, combining datasets across sites leads to two challenges: first, an increase in undesirable non-biological variance due to scanner and acquisition differences - the *harmonisation problem* - and second, data privacy concerns due to the inherently personal nature of medical imaging data, meaning that sharing them across sites may risk violation of privacy laws. To overcome these restrictions, we propose **FedHarmony**: a harmonisation framework operating in the federated learning paradigm. We show that to remove the scanner-specific effects, for our scenario we only need to share the mean and standard deviation of the learned features, helping to protect individual subjects' privacy. We demonstrate our approach across a range of realistic data scenarios, using real multi-site data from the ABIDE dataset, thus showing the potential utility of our method for MRI harmonisation across studies. Our code is available at <https://github.com/nkdinsdale/FedHarmony>.

Keywords: Harmonisation · Federated Learning · Domain Adaptation.

1 Introduction

Although some large scale projects, such as the UK Biobank, exist, the majority of neuroimaging datasets remain small. This necessitates the combination of data from multiple sites and scanners, both for statistical power and to represent the breadth of biological variability. However, combining data across scanners leads to an increase in non-biological variance, due to differences such as acquisition protocol and hardware [9,10,24]. Thus, we need *harmonisation methods* to enable joint unbiased analysis of data from different scanners and studies.

However, data cannot simply be pooled (combined) across imaging sites without ensuring compliance with data privacy laws, particularly if we wish to use representative clinical imaging data, as the sharing of medical images is covered

by legislation such as GDPR [2] and HIPAA [13]. Federated learning (FL) has been proposed as a method to train models on distributed data [11], where data are kept on their local servers, the users train local models on their own private datasets and then share only the weights or gradients of the trained models. FL has the potential to become the standard paradigm for multisite imaging, with early studies demonstrating its feasibility in neuroimaging [18].

Direct translation of existing harmonisation methods into FL frameworks is non-trivial. Most deep learning methods for harmonisation are based on generative frameworks [3,12,26,27], and, although federated equivalents to GANs and VAEs are being developed [16,25], additional challenges exist for harmonisation approaches that require simultaneous access to source and target data [8]. Additionally, many methods require paired data – not possible with distributed data and unlikely to exist in large multisite studies. Further, most generative methods are data-hungry [3,27], which casts into doubt whether sufficient data would be available at local sites.

Alternative methods for harmonisation frame the task as a domain adaptation problem (DA) [5], as the goal mirrors the harmonisation problem: removal of information regarding domain (scanner or acquisition protocol) while retaining the true variance of interest (the biological signal). In [5], the DA approach has been demonstrated for a range of tasks and network architectures. The harmonisation approach in [5] is built upon the DA framework proposed in [20] and has since been translated into the federated setting [14]. In [14] the domain adaptation across local sites is achieved by sharing feature embeddings for local data points in a global shared knowledge bank, such that a global domain predictor can be trained. However, sharing the features still leads to privacy concerns [8], as images are potentially recoverable from the features. Many of the other FL DA methods proposed cannot be adapted for the harmonisation problem, as, for instance, they produce domain-specific models or ensembles [8,15,21,23], where the final predictions depend on the source of the data.

Therefore, we present **FedHarmony**: a method to adapt the harmonisation framework proposed in [5] for distributed data, while minimising shared information. We aim to train a model that performs equally well on the task of interest across all of the imaging sites while using features that are invariant to the imaging site, through a horizontal federated framework [22]. We show that by approximating the learned feature embeddings as Gaussian distributions, we can share only a mean and standard deviation per feature which, by definition, contain no identifying information for individual subjects. Alongside standard privacy-protecting approaches [7,19], sharing only these summary statistics would make the approach robust to honest-but-curious adversaries [22]. By demonstrating our approach, **FedHarmony**, for a range of data settings, we show its viability for data harmonisation for multisite distributed imaging studies.

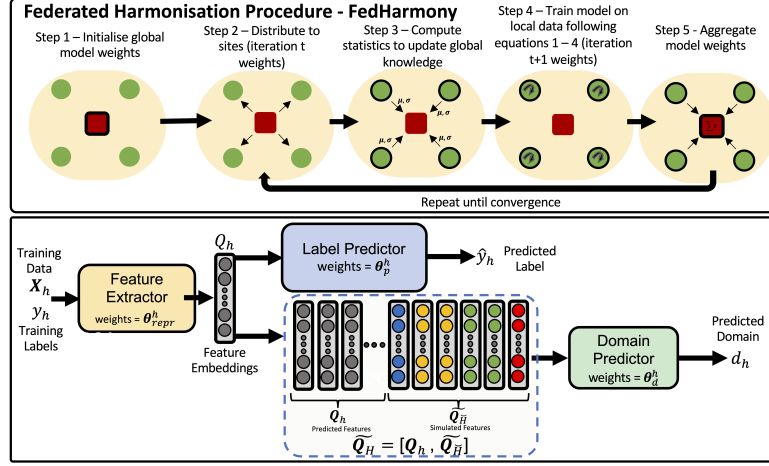


Fig. 1: FedHarmony Procedure and high level general network architecture.

2 Methods

Consider the scenario in which we have multisite MRI data, where the data for each site h are stored on their local server, for all sites H . For each site, we have pairs of training data $D_h = \{\mathbf{X}_h, \mathbf{y}_h\}$ where \mathbf{X}_h is the input data and \mathbf{y}_h is the label, and we consider each site to be a separate data domain. We wish to create a model that performs as well as possible on each imaging site, whilst having predictions that are invariant to the acquisition site. A global model is constructed, which following [5] is formed of a feature extractor (with weights θ_{repr}), a label predictor (θ_p) and a domain predictor (θ_d) (Fig. 1). This model is initialised randomly (step 1), and then the **FedHarmony** procedure begins, as shown in Fig. 1, with the global model weights for the three network components, $\theta_{repr,t}$, $\theta_{p,t}$ and $\theta_{d,t}$, being sent to the local sites; t indicates these are the current weights for iteration t and thus each site receives the same initialisation (step 2).

Update Global Knowledge Store (Step 3): To remove scanner information adversarially, the different sites' nodes need to have an understanding of the various scanner characteristics. Although there is precedent for achieving this through sharing the feature embeddings from local sites to the global site [14], we aim to minimise the quantity of shared information to help protect individual privacy. It has been shown that domain-specific characteristics can be encoded to enable contrastive learning [6], by sharing the mean and standard deviation of the feature embeddings following a BoxCox transformation [1] (an invertible transform which aims to make the distribution maximally Gaussian), suggesting that this might be sufficient information to also remove scanner information. However, in many multisite MRI studies, small amounts of local data are available that

may be insufficient to fit a BoxCox transformation. Thus, we simplify the transformation by assuming that the features will already be normally distributed and that we can characterise $\mathbf{Q}_h = \text{featureExtractor}(\mathbf{X}_h, \boldsymbol{\theta}_{repr,t})$ by their mean and standard deviations. Therefore the total information shared per site is a mean and standard deviation per feature, which inherently has no information about individuals, and we collate a global information store $\{\boldsymbol{\mu}_h, \boldsymbol{\sigma}_h : h \in H\}$.

Training Procedure (Step 4): Local training is then controlled by three loss functions, based on [5]. First is the main task loss L_p , in which the feature extractor and the label predictor are updated. If $\Theta_t = [\boldsymbol{\theta}_{repr,t}, \boldsymbol{\theta}_{p,t}]$ are the current global weights (iteration t) for the feature extractor and the label predictor, and $\Theta_{t+1}^h = [\boldsymbol{\theta}_{repr,t+1}^h, \boldsymbol{\theta}_{p,t+1}^h]$ are the updated $(t+1)$ weights for site h , then:

$$L_p(\mathbf{X}_h, \mathbf{y}_h; \boldsymbol{\theta}_{repr,t+1}^h, \boldsymbol{\theta}_{p,t+1}^h) = \frac{1}{N_h} \sum_{j=1}^{N_h} L_{task}(\mathbf{y}_{h,j}, \hat{\mathbf{y}}_{h,j}) + \mu L_{prox}(\Theta_t, \Theta_{t+1}^h) \quad (1)$$

where L_{task} is the task-specific loss function, averaged over the training samples available for the given site, N_h , and $\mathbf{y}_{h,j}$ and $\hat{\mathbf{y}}_{h,j}$ are the true and predicted labels for the j^{th} example for the h^{th} imaging site. L_{prox} is a proximal loss term based on FedProx [17] that penalises weight deviations from the global model that do not sufficiently improve the main task loss and has been shown to aid FL when the data is non-iid [17], as is inherently true for multisite MRI data, hence the harmonisation problem. Thus, if μ is a constant, weighting the contributions from the two losses, then:

$$L_{prox}(\Theta_t, \Theta_{t+1}^h) = \|\boldsymbol{\theta}_{repr,t} - \boldsymbol{\theta}_{repr,t+1}^h\|^2 + \|\boldsymbol{\theta}_{p,t} - \boldsymbol{\theta}_{p,t+1}^h\|^2 \quad (2)$$

The second loss function, L_d , then updates the domain predictor to be able to discriminate between the sites. Data is only available for the local site h and thus can only directly generate the features, \mathbf{Q}_h , needed to train the domain predictor for this local site. For all the other sites \check{H} we generate example features randomly using the shared knowledge store; thus, for site \check{h} in $\check{H} : \mathbf{Q}_{\check{h}} \sim \mathcal{N}(\boldsymbol{\mu}_{\check{h}}, \boldsymbol{\sigma}_{\check{h}})$. The generated features, $\tilde{\mathbf{Q}}_{\check{H}} = \{\tilde{\mathbf{Q}}_{\check{h}} \forall \check{h} \in \check{H}\}$ are then concatenated with the true features, \mathbf{Q}_h , such that $\tilde{\mathbf{Q}}_H^h = [\mathbf{Q}_h, \tilde{\mathbf{Q}}_{\check{H}}]$. The number of simulated subjects generated for each site \check{h} is chosen randomly such that the final batchsize is twice that of the original; the features and corresponding domain (site) labels are then shuffled. We then update the parameters of the domain predictor, using categorical cross-entropy to assess the site information remaining in \mathbf{Q}_h :

$$L_d(\tilde{\mathbf{Q}}_H^h, \mathbf{d}; \boldsymbol{\theta}_{d,t+1}^h) = - \sum_{h=1}^H \mathbb{1}[d = h] \log(p_h) \quad (3)$$

where p_h are the softmax outputs of the domain classifier. The third loss is the confusion loss L_{conf} , which removes the site information, by penalising deviation in p_h from a uniform distribution:

$$L_{conf}(\mathbf{X}_h, \boldsymbol{\theta}_{d,t+1}^h; \boldsymbol{\theta}_{repr,t+1}^h) = - \sum_{h=1}^H \frac{1}{H} \log(p_h) \quad (4)$$

Therefore, the overall method can be considered to minimise the total loss function $L = L_p + \alpha L_d + \beta L_{conf}$ where α and β control the relative contributions of the different loss functions. The domain loss (Eq. 3) and the confusion loss (Eq. 4) act in opposition to each other, and so must be updated iteratively, with three iterations per epoch. The procedure is iterated through for E local epochs.

Weight Aggregation (Step 5): After each site has completed the E local epochs, the local weights are returned to the global node to be aggregated. In [5] it was shown that the harmonisation process was aided through evaluating each site separately in the loss function rather than averaging over all data points. Given that FedAvg [11] is equivalent to averaging over all training subjects for FL, we alter the aggregation to have equal contributions from each site: $\theta_{t+1} \leftarrow \frac{1}{H} \sum_{h=1}^H \theta_{t+1}^h$ where the aggregation is completed separately for θ_{repr} , θ_p and θ_d . We term this aggregation step FedEqual. Once the aggregation is complete, the training loops through stages 2 - 5 (Fig. 1) until convergence.

Semi-supervised setting: We then explore the scenario where labels are only available for a subset of sites, which is realistic, as labels are expensive to generate in terms of time and expertise. Thus, we also consider the semi-supervised setting, adjusting the framework such that local training is only completed for sites with labels available, but we continue to update the global knowledge store for all sites. Therefore, the removal of the site information still considers all sites.

3 Implementation Details

For our experiments we use data from the ABIDE dataset⁵ [4], for the task of age prediction, using T1 MR images as \mathbf{X} and the ages as the labels \mathbf{y} . The MR images from each site were processed using FSL anat⁶. Subjects were rejected where the pipeline failed or where data was missing. Four sites (Trinity, NYU, UCLA, Yale) were chosen for our experiments, so as to span both age distributions and subject numbers (Fig. 2). The data were split into training/validation/test sets as 70%/10%/20%, yielding a maximum of 127 subjects for training (NYU) and a minimum of 35 (Trinity). Our networks and baselines were all trained with 5-fold cross-validation until no improvement in the average mean absolute error (MAE) was reported on the validation data for all local sites. We present the results on the held-out testing data in all cases.

For the network, we used the VGG-based architecture used in [5]⁷. The images were resized to (128, 240, 160) and normalised to have zero mean and unit standard deviation. As demonstrated by [5], the framework is flexible and should be applicable to many feedforward architectures and tasks, with age prediction being used to demonstrate the approach due to the availability of labels across many sites. The implementation was in Python 3.6.8 and PyTorch 1.10.2. The

⁵ Data from: https://fcon_1000.projects.nitrc.org/indi/abide/

⁶ https://fsl.fmrib.ox.ac.uk/fsl/fslwiki/fsl_anat

⁷ code from: https://github.com/nkdinsdale/Unlearning_for_MRI_harmonisation

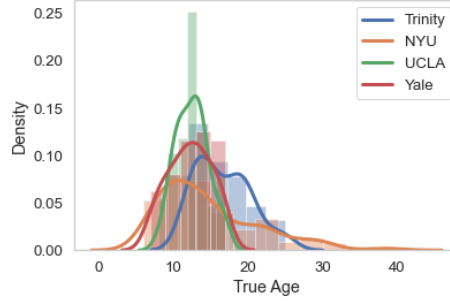


Fig. 2: Normalised age distributions for the 4 sites from the ABIDE dataset: Trinity: 49 subjects, 16.7 ± 3.6 years (range: 12-25), NYU: 182 subjects, 14.7 ± 6.6 (6-39), UCLA: 99 subjects, 12.5 ± 2.2 (8-17), Yale: 56 subjects, 12.2 ± 2.8 (7-17).

networks were trained on a 16GB P100 GPU, taking up to 5 minutes per local epoch (NYU), with a batchsize of 16, and requiring 50-75 rounds of communication to converge (**FedHarmony** required 61 rounds and **FedAvg** required 56). The following hyperparameters were used for all experiments based on [5,17], with the same values applied for each local site: batchsize=16, learning rate= 1×10^{-4} with Adam optimiser, local epochs (E)=10, $\mu=0.01$, $\alpha=1$, and $\beta=100$.

4 Results & Discussion

Fully Supervised: We first compared our method to several baselines: training on NYU data only (the largest site, representing where data cannot be shared and we have no federated learning framework), training on all sites normally (centralised data), **FedAvg** [11], **FedProx** [17] and our proposed **FedEqual** aggregation method. We also completed an ablation study, exploring the effect of the aggregation and proximal losses alongside the harmonisation. To compare methods, we evaluated both the mean absolute error (MAE), where we want the best performance possible across all sites, and the scanner classification accuracy (SCA): the performance of a domain classifier trained on \mathbf{Q} at convergence, where the aim is to achieve random chance (25%). Note that the domain predictor was able to achieve 95% on the feature embeddings when the age prediction was trained with NYU only, and so is clearly able to identify scanner information.

Considering the results in Table 1, we were able to train federated models (**FedAvg**) to perform as well as standard centralised training. **FedHarmony** led to improvement both in terms of average performance across all sites and in the reduction of SCA to close to random chance. This improvement indicates the success of the DA approach, as information from across sites is being used to aid predictions even though only summary statistics of the features are being shared. The ablation study shows the need for the three additions to the **FedAvg** baseline, with the L_{prox} (Eq. 2) being vital for both stability of training and the ability to remove scanner information. Fig. 3 shows a PCA of \mathbf{Q} comparing

FedAvg and FedHarmony, where it is evident that the harmonisation increases the feature overlap of the different sites.

Table 1: **Fully supervised results:** Component represents additions relative to FedAvg [11]: Equal = Aggregation as FedEqual; Prox = Proximal Loss as Eq. 2; Harm = Harmonisation, otherwise local training only considers the main task loss. SCA = Scanner Classification Accuracy of domain classifier retrained on \mathbf{Q} at the end of training, where random chance (25%) is the goal. NYU only and all sites = centralised training. FedHarmony is our proposed approach. * = significant improvement over next performing method (paired t-test, $p < 0.001$).

Method	Component			Site MAE				Average MAE	SCA (%)
	Equal	Prox	Harm	NYU n=40	Yale n=14	UCLA n=20	Trinity n=10		
NYU Only				5.26 \pm 4.37	2.38 \pm 1.42	3.22 \pm 2.26	9.66 \pm 7.89	5.13	95
All Sites				5.10 \pm 4.56	2.32 \pm 1.50	2.87 \pm 2.13	4.53 \pm 3.49	3.70	86
FedAvg [11]				5.26 \pm 4.74	1.99 \pm 1.21*	2.57 \pm 1.88	4.94 \pm 3.80	3.69	64
FedProx [17]		✓		5.21 \pm 4.70	2.15 \pm 1.44	2.61 \pm 1.73	4.47 \pm 3.56	3.61	62
FedEqual	✓			6.07 \pm 3.99	2.06 \pm 1.31	2.61 \pm 2.30	5.43 \pm 3.00	4.06	42
Ablation A			✓	12.21 \pm 10.31	4.20 \pm 3.61	6.21 \pm 3.21	10.32 \pm 8.31	8.23	95
Ablation B	✓		✓	5.27 \pm 4.90	2.27 \pm 1.13	2.66 \pm 1.98	4.81 \pm 3.98	3.75	34
Ablation C		✓	✓	4.93 \pm 4.31*	2.13 \pm 1.91	2.26 \pm 1.15	4.36 \pm 3.37	3.42	35
FedHarmony	✓	✓	✓	5.14 \pm 4.52	2.04 \pm 1.25	2.13 \pm 1.45*	4.34 \pm 3.65	3.42	29

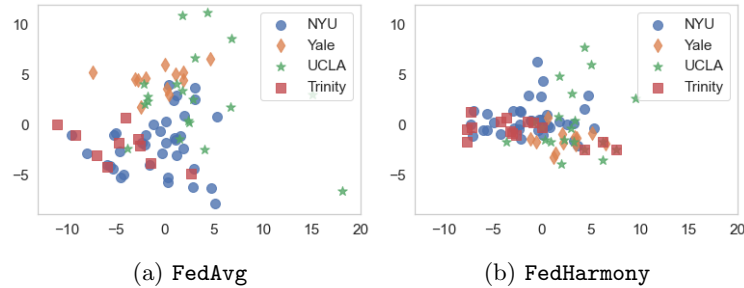


Fig. 3: PCA of \mathbf{Q} for FedAvg and FedHarmony, showing, as expected, increased overlap across sites with FedHarmony.

Semi-supervised: We then simulated two semi-supervised scenarios, where training labels were only available for a subset of the sites. First we considered when only one site had available training labels, NYU, as a multitarget DA task, which could for example represent when the organising site has labelled data. In Table 2 we compare to standard training (FedAvg and FedProx are equivalent to standard training for a single labelled site). Using FedHarmony, we were able to train the network to perform as well as fully supervised centralised training,

Table 2: **Semi-supervised results:** *NYU only* = centralised training only NYU; *No Trinity* = centralised training no Trinity; *All sites* = centralised training. SCA = Scanner Classification Accuracy of domain classifier retrained on \mathbf{Q} at the end of training, where random chance (25%). **FedHarmony** was trained locally for only labelled sites, and global shared information used for all sites. * = significant improvement over next performing method (paired t-test, $p < 0.001$).

Method	Site MAE				Average MAE	SCA (%)
	NYU n=40	Yale n=14	UCLA n=20	Trinity n=10		
a) Semi-supervised 1 Site						
NYU Only	5.26 ± 4.37	3.22±2.26	3.22±2.26	9.66±7.89	5.13	95
FedHarmony	5.26 ± 4.15	2.20 ± 1.79*	2.91 ± 1.56*	4.09 ± 3.32*	3.61*	30
b) Semi-supervised 3 Site						
No Trinity	5.36±4.65	1.81 ± 1.66*	1.83 ± 1.16*	7.29±5.00	4.07	96
FedProx [17]	5.15±4.76	2.15±1.38	2.62±1.82	5.96±4.05	3.96	74
FedHarmony	5.12 ± 4.60*	2.39±1.67	2.16±1.40	4.69 ± 4.19*	3.59*	26
All Sites	5.10±4.56	2.32±1.50	2.87±2.13	4.53±3.49	3.70	86

while removing the scanner information, showing the power of the shared global knowledge. We also considered when three sites have labels and one does not (Trinity), for example, if a new site joins the study. **FedHarmony** led to improvement in performance again, even compared to having labels for all sites, while removing the scanner information. These results show the suitability of our approach across realistic data scenarios for multisite MRI studies.

Suitability of Gaussian Fit: In Fig. 4 we plot the standard deviation of 100 mean estimates for \mathbf{Q}_{NYU} from both the direct Gaussian fit and the BoxCox transformation for increasing numbers of samples. It is evident that the direct Gaussian fit leads to more consistent estimates, indicating increased suitability in low data regimes compared to the BoxCox transform, as the representation on the training data is more likely to also represent the testing data.

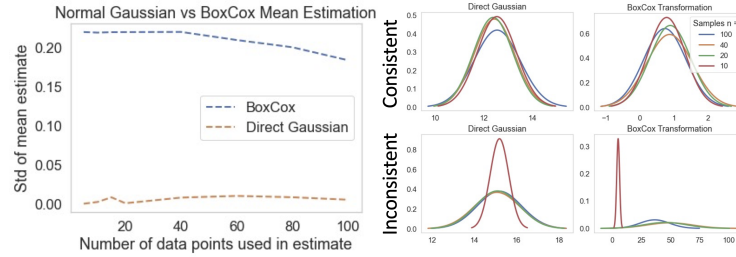


Fig. 4: We plot the standard deviation of 100 mean estimates of the fit for \mathbf{Q}_{NYU} , for increasing numbers of random samples from the full \mathbf{Q}_{NYU} , alongside consistent and inconsistent fits. Note: the estimated means will differ, as the BoxCox transforms the data to be maximally Gaussian rather than fitting the data.

5 Conclusion

We have presented **FedHarmony**, a method to allow harmonisation of MRI data in a federated learning scenario, through development of a domain adaptation approach with minimal information sharing, outperforming baseline and FL approaches across realistic data scenarios. Alongside standard privacy-protecting approaches, such as differential privacy, our approach would enable training models on multisite MRI data while maintaining individual subjects' privacy. Future work will explore application to different neuroimaging tasks.

6 Acknowledgements

ND is supported by a Academy of Medical Sciences Springboard Award. MJ is supported by the National Institute for Health Research, Oxford Biomedical Research Centre, and this research was funded by the Wellcome Trust [215573/Z/19/Z]. WIN is supported by core funding from the Wellcome Trust [203139/Z/16/Z]. AN is grateful for support from the UK Royal Academy of Engineering under the Engineering for Development Research Fellowships scheme and the Academy of Medical Sciences. The computational aspects of this research were supported by the Wellcome Trust Core Award [Grant Number 203141/Z/16/Z] and the NIHR Oxford BRC. The views expressed are those of the author(s) and not necessarily those of the NHS, the NIHR or the Department of Health.

References

1. Box, G.E.P., Cox, D.R.: An analysis of transformations. *Journal of the Royal Statistical Society. Series B (Methodological)* **26**(2), 211–252 (1964). <https://doi.org/10.2307/2984418>
2. Consulting, I.: General Data Protection Regulation (GDPR) (Sep 2019), <https://gdpr-info.eu/>
3. Dewey, B., Zhao, C., Reinhold, J., Carass, A., Fitzgerald, K., Sotirchos, E., Saidha, S., Oh, J., Pham, D., Calabresi, P., Zijl, P., Prince, J.: DeepHarmony: A deep learning approach to contrast harmonization across scanner changes. *Magnetic Resonance Imaging* **64** (07 2019)
4. Di Martino, A., Yan, C.G., Li, Q., Denio, E., Castellanos, F., Alaerts, K., Anderson, J., Assaf, M., Bookheimer, S., Dapretto, M., Deen, B., Delmonte, S., Dinstein, I., Birgit, E.W., Fair, D., Gallagher, L., Kennedy, D., Keown, C., Keyzers, C., Milham, M.: The Autism brain imaging data exchange: Towards large-scale evaluation of the intrinsic brain architecture in Autism. *Molecular psychiatry* **19** (06 2013). <https://doi.org/10.1038/mp.2013.78>
5. Dinsdale, N.K., Jenkinson, M., Namburete, A.I.: Deep learning-based unlearning of dataset bias for MRI harmonisation and confound removal. *NeuroImage* **228**, 117689 (2021). <https://doi.org/https://doi.org/10.1016/j.neuroimage.2020.117689>

6. Dong, N., Voiculescu, I.: Federated contrastive learning for decentralized unlabeled medical images. In: Medical Image Computing and Computer Assisted Intervention – MICCAI 2021 (09 2021)
7. Dwork, C.: Differential privacy. In: Bugliesi, M., Preneel, B., Sassone, V., Wegener, I. (eds.) Automata, Languages and Programming. pp. 1–12 (2006)
8. Feng, H., You, Z., Chen, M., Zhang, T., Zhu, M., Wu, F., Wu, C., Chen, W.: Kd3a: Unsupervised multi-source decentralized domain adaptation via knowledge distillation. In: ICML (2021)
9. Han, X., Jovicich, J., Salat, D., Kouwe, A., Quinn, B., Czanner, S., Busa, E., Pacheco, J., Albert, M., Killiany, R., Maguire, R., Rosas, H., Makris, N., Dale, A., Dickerson, B., Fischl, B.: Reliability of MRI-derived measurements of human cerebral cortical thickness: The effects of field strength, scanner upgrade and manufacturer. *NeuroImage* **32**, 180–94 (09 2006). <https://doi.org/10.1016/j.neuroimage.2006.02.051>
10. Jovicich, J., Czanner, S., Greve, D., Haley, E., Kouwe, A., Gollub, R., Kennedy, D., Schmitt, F., Brown, G., Macfall, J., Fischl, B., Dale, A.: Reliability in multi-site structural MRI studies: Effects of gradient non-linearity correction on phantom and human data. *NeuroImage* **30**, 436–43 (05 2006). <https://doi.org/10.1016/j.neuroimage.2005.09.046>
11. McMahan, H.B., Moore, E., Ramage, D., Hampson, S., y Arcas, B.A.: Communication-efficient learning of deep networks from decentralized data. In: AISTATS (2017)
12. Moyer, D., Ver Steeg, G., Tax, C., Thompson, P.: Scanner invariant representations for diffusion MRI harmonization. *Magnetic Resonance in Medicine* (04 2020). <https://doi.org/10.1002/mrm.28243>
13. Office for Civil Rights US: The hipaa privacy rule (Jul 2021), <https://www.hhs.gov/hipaa/for-professionals/privacy/index.html>
14. Peng, X., Huang, Z., Zhu, Y., Saenko, K.: Federated adversarial domain adaptation. *ArXiv abs/1911.02054* (2020)
15. Peterson, D., Kanani, P.H., Marathe, V.J.: Private federated learning with domain adaptation. *ArXiv abs/1912.06733* (2019)
16. Rasouli, M., Sun, T., Rajagopal, R.: Fedgan: Federated generative adversarial networks for distributed data. *ArXiv* (06 2020)
17. Sahu, A., Li, T., Sanjabi, M., Zaheer, M., Talwalkar, A., Smith, V.: On the convergence of federated optimization in heterogeneous networks. *ArXiv* (12 2018)
18. Sheller, M.J., Reina, G.A., Edwards, B., Martin, J., Bakas, S.: 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. pp. 92–104 (2019). <https://doi.org/10.1007/978-3-030-11723-8-9>
19. Sweeney, L.: K-anonymity: A model for protecting privacy. *Int. J. Uncertain. Fuzziness Knowl.-Based Syst.* **10**(5), 557–570 (Oct 2002). <https://doi.org/10.1142/S0218488502001648>
20. Tzeng, E., Hoffman, J., Darrell, T., Saenko, K.: Simultaneous deep transfer across domains and tasks. In: ICCV (2015)
21. Yang, L., Beliard, C., Rossi, D.: Heterogeneous data-aware federated learning. *ArXiv* (11 2020)
22. Yang, Q., Liu, Y., Chen, T., Tong, Y.: Federated machine learning: Concept and applications. *ACM Transactions on Intelligent Systems and Technology* **10**, 1–19 (01 2019). <https://doi.org/10.1145/3298981>

23. Yao, C.H., Gong, B., Cui, Y., Qi, H., Zhu, Y., Yang, M.H.: Federated multi-target domain adaptation. ArXiv (2021)
24. Yu, M., Linn, K., Cook, P., Phillips, M., McInnis, M., Fava, M., Trivedi, M., Weissman, M., Shinohara, R., Sheline, Y.: Statistical harmonization corrects site effects in functional connectivity measurements from multisite fMRI data. *Human Brain Mapping* **39** (07 2018). <https://doi.org/10.1002/hbm.24241>
25. Zhang, K., Jiang, Y., Seversky, L., Xu, C., Liu, D., Song, H.: Federated variational learning for anomaly detection in multivariate time series. ArXiv (08 2021)
26. Zhao, F., Wu, Z., Wang, L., Lin, W., Xia, S., Li, G.: Harmonization of infant cortical thickness using surface-to-surface cycle-consistent adversarial networks. *Medical Image Computing and Computer Assisted Intervention - Conference Proceedings* pp. 475–483 (10 2019)
27. Zuo, L., Dewey, B.E., Liu, Y., He, Y., Newsome, S.D., Mowry, E.M., Resnick, S.M., Prince, J.L., Carass, A.: Unsupervised mr harmonization by learning disentangled representations using information bottleneck theory. *NeuroImage* **243**, 118569 (2021). <https://doi.org/https://doi.org/10.1016/j.neuroimage.2021.118569>