



HHS Public Access

Author manuscript

Med Image Comput Comput Assist Interv. Author manuscript; available in PMC 2023 March 22.

Published in final edited form as:

Med Image Comput Comput Assist Interv. 2021 ; 12903: 181–190. doi:10.1007/978-3-030-87199-4_17.

Harmonization with Flow-Based Causal Inference

Rongguang Wang^{1,2}, Pratik Chaudhari^{1,3}, Christos Davatzikos^{1,2,4}

¹Department of Electrical and Systems Engineering, University of Pennsylvania, Philadelphia, USA

²Center for Biomedical Image Computing and Analytics (CBICA), Philadelphia, USA

³General Robotics, Automation, Sensing and Perception Laboratory (GRASP), Philadelphia, USA

⁴Department of Radiology, Perelman School of Medicine, University of Pennsylvania, Philadelphia, USA

Abstract

Heterogeneity in medical data, e.g., from data collected at different sites and with different protocols in a clinical study, is a fundamental hurdle for accurate prediction using machine learning models, as such models often fail to generalize well. This paper leverages a recently proposed normalizing-flow-based method to perform counterfactual inference upon a structural causal model (SCM), in order to achieve harmonization of such data. A causal model is used to model observed effects (brain magnetic resonance imaging data) that result from known confounders (site, gender and age) and exogenous noise variables. Our formulation exploits the bijection induced by flow for the purpose of harmonization. We infer the posterior of exogenous variables, intervene on observations, and draw samples from the resultant SCM to obtain counterfactuals. This approach is evaluated extensively on multiple, large, real-world medical datasets and displayed better cross-domain generalization compared to state-of-the-art algorithms. Further experiments that evaluate the quality of confounder-independent data generated by our model using regression and classification tasks are provided.

Keywords

Harmonization; Causal inference; Normalizing flows

1 Introduction

Deep learning models have shown great promise in medical imaging diagnostics [11] and predictive modeling with applications ranging from segmentation tasks [19] to more complex decision-support functions for phenotyping brain diseases and personalized prognosis. However deep learning models tend to have poor reproducibility across hospitals, scanners, and patient cohorts; these high-dimensional models tend to overfit to specific

rgw@seas.upenn.edu .

Electronic supplementary material The online version of this chapter (https://doi.org/10.1007/978-3-030-87199-4_17) contains supplementary material, which is available to authorized users.

datasets and generalize poorly across training data [6]. One potential solution to the above problem is to train on very large and diverse databases but this can be prohibitive, because data may change frequently (e.g., new imaging devices are introduced) and gathering training labels for medical images is expensive. More importantly, even if it were possible to train a model on data that covers all possible variations across images, such a model would almost certainly sacrifice accuracy in favor of generalization—it would rely on coarse imaging features that are stable across, say imaging devices and patient populations, and might fail to capture more subtle and informative detail. Methods that can tackle heterogeneity in medical data without sacrificing predictive accuracy are needed, including methods for “data harmonization”, which would allow training a classifier on, say data from one site, and obtaining similar predictive accuracy on data from another site.

Contributions.

We build upon a recently proposed framework [26] for causal inference, by modeling brain imaging data and clinical variables via a causal graph and focus on how causes (site, gender and age) result in the effects, namely imaging measurements (herein we use region of interest (ROI) volumes obtained by preprocessing brain MRI data). This framework uses a normalizing flow parameterized by deep networks to learn the structural assignments in a causal graph. We demonstrate how harmonization of data can be performed efficiently using counterfactual inference on such a flow-based causal model. Given a dataset pertaining to one site (source), we perform a counterfactual query to synthesize the dataset, as if it were from another site (target). Essentially, this amounts to the counterfactual question “what would the scans look like if they had been acquired from the same site”. We demonstrate results of such harmonization on regression (age prediction) and classification (predicting Alzheimer’s disease) tasks using several large-scale brain imaging datasets. We demonstrate substantial improvement over competitive baselines on these tasks.

2 Related Work

A wide variety of recent advances have been made to remove undesired confounders for imaging data, e.g., pertaining to sites or scanners [3,16,21,22,30,32]. Methods like ComBat [16,30], based on parametric empirical Bayes [20], produce site-removed image features by performing location (mean) and scale (variance) adjustments to the data. A linear model estimates location and scale differences in images features across sites while preserving confounders such as sex and age. In this approach, other unknown variations such as race and disease are removed together with the site variable, which might lead to inaccurate predictions for disease diagnosis. Generative deep learning models such as variational autoencoders (VAEs) [18] and generative adversarial networks (GANs) [12] have been used in many works [3,21,22,32]. These methods typically minimize the mutual information between the site variable and image embedding in the latent space, and learn a site-disentangled representation which can be used to reconstruct images from a different site. Unsupervised image-to-image translation has been used to map scans either between two sites [32] or to a reference domain [3] using models like CycleGAN [37]. Generative models are however challenging to use in practice: VAEs typically suffer from blurry reconstructions while GANs can suffer from mode collapse and convergence

issues. These issues are exacerbated for 3D images. In this paper, we focus on regions of interest (ROI) features. We extend the deep structural causal model of [26] which enables tractable counterfactual inference from single-site healthy MR images to multi-center pathology-associated scans for data harmonization. Besides qualitative examination of the counterfactuals performed in [26], we provide extensive quantitative evaluations and compare it with state-of-the-art harmonization baselines.

3 Method

Our method builds upon the causal inference mechanism proposed by Judea Pearl [27] and method of Pawlowski et al. [26] that allows performing counterfactual queries upon causal models parameterized by deep networks. We first introduce preliminaries of our method, namely, structural causal models, counterfactual inference, and normalizing flows, and then describe the proposed harmonization algorithm.

3.1 Building Blocks

Structural Causal Models (SCMs).—are analogues of directed probabilistic graphical models for causal inference [29,33]. Parent-child relationships in an SCM denote the effect (child) of direct causes (parents) while they only denote conditional independencies in a graphical model. Consider a collection of random variables $x = (x_1, \dots, x_m)$, an SCM given by $M = (S, P_e)$ consists of a collection $S = (f_1, \dots, f_m)$ of assignments $x_k = f_k(\epsilon_k; \text{pa}_k)$ where pa_k denotes the set of parents (direct causes) of x_k and noise variables ϵ_k are unknown and unmodeled sources of variation for x_k . Each variable x_k is independent of its non-effects given its direct causes (known as the causal Markov condition), we can write the joint distribution of an SCM as $P_M(x) = \prod_{k=1}^m P(x_k | \text{pa}_k)$; each conditional distribution here is determined by the corresponding structural assignment f_k and noise distribution [27]. Exogenous noise variables are assumed to have a joint distribution $P_e = \prod_{k=1}^m P(\epsilon_k)$, this will be useful in the sequel.

Counterfactual Inference.—Given a SCM, a counterfactual query is formulated as a three-step process: abduction, action, and prediction [27–29]. First, we predict exogenous noise ϵ based on observations to get the posterior $P_M(\epsilon | x) = \prod_{k=1}^m P_M(\epsilon_k | x_k, \text{pa}_k)$. Then comes intervention denoted by $\text{do}(\tilde{x}_k)$, where we replace structural assignments of variable x_k . Intervention makes the effect x_k independent of both its causes pa_k and noise ϵ_k and this results in a modified SCM $\tilde{M} = M_{\text{do}(\tilde{x}_k)} \equiv (\tilde{S}, P_M(\epsilon | x))$. Note that the noise distribution has also been modified, it is now the posterior $P_M(\epsilon | x)$ obtained in the abduction step. The third step, namely prediction involves predicting counterfactuals by sampling from the distribution $P_{\tilde{M}}(x)$ entailed by the modified SCM.

Learning a Normalizing Flow-Based SCM.—Given the structure of the SCM, learning the model involves learning the structure assignments S from data. Following [26], we next do so using normalizing flows parameterized by deep networks. Normalizing flows model a complex probability density as the result of a transformation applied to some simple probability density [7,10,23,24]; these transformations are learned using samples from the target. Given observed variables x and base density $\epsilon \sim \mathcal{P}(\epsilon)$, this involves finding an

invertible and differentiable transformation $x = f(\epsilon)$. The density of x is given by $p(x) = p(\epsilon) |\det \nabla f(\epsilon)|^{-1}$ where $\epsilon = f^{-1}(x)$ and $\nabla f(\epsilon)$ is the Jacobian of the flow $f: \epsilon \mapsto x$. The density $p(\epsilon)$ is typically chosen to be a Gaussian. Given a dataset $D = \{x^i \sim p(x)\}_{i=1}^n$ with n samples, a θ -parameterized normalizing flow f_θ can be fitted using a maximum-likelihood objective to obtain

$$\theta^* = \operatorname{argmax}_\theta \frac{1}{n} \sum_{i=1}^n \log p(\epsilon^i) - \log |\det \nabla f_\theta(\epsilon^i)|.$$

Here $\epsilon^i = f_\theta^{-1}(x^i)$. Parameterizing the normalizing flow using a deep network leads to powerful density estimation methods. This approach can be easily extended to conditional densities of the form $p(x_k | \text{pa}_k)$ in our SCM.

3.2 Harmonization Using Counterfactual Inference in a Flow-Based SCM

Given the structure of a SCM, we fit conditional flows $f_{\theta_k}: \epsilon_k \mapsto x_k$ that map exogenous noise to effect x_k given parents pa_k for all nodes in the SCM. We will denote the combined flow for all nodes in the SCM as f_θ which maps noise $\epsilon^i = (\epsilon^i_1, \dots, \epsilon^i_m)$ to observations $x^i = (x^i_1, \dots, x^i_m)$ in the dataset; the corresponding SCM is denoted by M_θ . Focus on a particular datum x^i in the dataset. The abduction step simply computes $\epsilon^i = f_\theta^{-1}(x^i)$. Formally this corresponds to computing the posterior distribution $P_{M_\theta}(\epsilon | x^i)$. Intervention uses the fact that the flow models a conditional distribution and replaces (intervenes) the value of a particular variable, say $x^i_k \leftarrow \tilde{x}^i_k$; this corresponds to the operation $\text{do}(\tilde{x}_k)$. The variable x_k is decoupled from its parents and exogenous noise which corresponds to a modified structural assignment \tilde{f}_{θ_k} and results in a new SCM \tilde{M}_θ . We can now run the same flow \tilde{f}_θ forwards using samples ϵ^i from the abduction step to get samples from $P_{\tilde{M}_\theta}(x)$ which are the counterfactuals. Figure 1 shows an example SCM for brain imaging data and shows we perform counterfactual queries to remove site effects.

4 Experimental Results

4.1 Setup

Datasets.—We use 6,921 3D T1-weighted brain magnetic resonance imaging (MRI) scans acquired from multiple scanners or sites in Alzheimer’s Disease Neuroimaging Initiative (ADNI) [15] and the iSTAGING consortium [13] which consists of Baltimore Longitudinal Study of Aging (BLSA) [2,31], Study of Health in Pomerania (SHIP) [14] and the UK Biobank (UKBB) [34]. Detailed demographic information of the datasets is provided in the supplementary material. We first perform a sequence of preprocessing steps on these images, including bias-field correction [35], brain tissue extraction via skull-stripping [8], and multi-atlas segmentation [9]. Each scan is then segmented into 145 anatomical regions of interests (ROIs) spanning the entire brain, and finally volumes of the ROIs are taken as the features. We first perform age prediction task using data from the iSTAGING consortium for participants between ages 21–93 years. We then demonstrate our method for classification of Alzheimer’s disease (AD) using the ADNI dataset where the diagnosis

groups are cognitive normal (CN) and AD; this is a more challenging problem than age prediction.

Implementation.—We implement flow-based SCM with three different flows (affine, linear and quadratic autoregressive splines [7,10]) using PyTorch [25] and Pyro [4]. We use a categorical distribution for sex and site, and real-valued normalizing flow for other structural assignments. A linear flow and a conditional flow (conditioned on activations of a fully-connected network that takes age, sex and scanner ID as input) are used as structural assignments for age and ROI features respectively. The density of exogenous noise is standard Gaussian. For training, we use Adam [17] with batch-size of 64, initial learning rate 3×10^{-4} and weight decay 10^{-4} . We use a staircase learning rate schedule with decay milestones at 50% and 75% of the training duration. All models are trained for at least 100 epochs. Implementation details for the SCM and the classifier, and the best validation log-likelihood for each model are shown in the supplementary material.

Baselines.—We compare with a number of state-of-the-art algorithms: invariant risk minimization (IRM) [1], ComBat [16,30], ComBat++ [36], and CovBat [5] on age regression and Alzheimer’s disease classification. IRM learns an invariant representation that the optimal predictor using this representation is simultaneously optimal in all environments. We implement IRM and ComBat algorithms with publicly available code from the original authors. We also show results obtained by training directly on the target data which acts as upper-bound on the accuracy of our harmonization.

4.2 Evaluation of the Learned Flow-Based SCM

We explore three normalizing flow models: affine, linear autoregressive spline [7], and quadratic autoregressive spline [10]. Implementation details and their log-likelihood are in the supplementary material. For both iSTAGING and ADNI datasets, the log-likelihood improves consistently with the model’s expressive power. Spline-based autoregressive flow models (17.22 log-likelihood for linear-spline and 17.24 for quadratic-spline) are better for density estimation than an affine flow model (1.88 log-likelihood). A quadratic spline-based model obtains slightly higher log-likelihood than the linear-spline model on the iSTAGING dataset.

We next show the feature (hippocampus volume) distributions of raw data, ComBat [16,30] transformed data and the data generated from the flow-based SCM in Fig. 2. We find that the feature distributions of ComBat are not consistent with those of raw data; ComBat transformed feature distributions show similar means (all shifted to an average value across sites) which arises from removing site-dependent location and scale effects. The third panel shows data generated from counterfactual queries with the flow-based SCM fitted on BLSA-3T. Feature distributions for the SCM-data are similar to those of the raw data. This can be attributed to the fact that our method preserves the unknown confounders (subject-specific information due to biological variability, such as race, gene, and pathology AD/CN) by capturing them as exogenous noise in the SCM.

4.3 Age Prediction

In Table 1, we compare the mean average error (MAE) of age prediction for a regression trained on raw data, site-removed data generated by ComBat [16,30] and its variants [5,36], IRM [1] and counterfactuals generated by our flow-based SCM. All models are trained on BLSA-3T (source site) and then tested on BLSA-1.5T, UKBB, and SHIP separately. We find that model (SrcOnly) trained on the source site with raw data cannot generalize to data from the target site. Models trained with site-removed data generated by ComBat generalize much better compared ones trained on raw data (SrcOnly), whereas IRM shows marginal improvement compared to SrcOnly. All variants (affine, linear-spline, quadratic-spline) of flow-based SCM show substantially smaller MAE; quadratic-spline SCM outperforms the other methods on all target sites.

4.4 Classification of Alzheimer's Disease

In Table 2, we show the accuracy of a classifier trained on raw data, Combat-harmonized data and SCM-generated counterfactuals for the ADNI dataset; this is a binary classification task with classes being CN (cognitive normal) and AD. All classifiers are trained on source sites (ADNI-1 or ADNI-2) and evaluated on target sites (ADNI-2 or ADNI-1 respectively). The classifier works poorly on the target site without any harmonization (SrcOnly). ComBat-based methods show a smaller gap between the accuracy on the source and target site; IRM improves upon this gap considerably. Harmonization using our flow-based SCM, in particular the Q-Spline variant, typically achieves higher accuracies on the target site compared to these methods.

5 Conclusion

This paper tackles harmonization of data from different sources using a method inspired from the literature on causal inference [26]. The main idea is to explicitly model the causal relationship of known confounders such as sex, age, and site, and ROI features in a SCM that uses normalizing flows to model probability distributions. Counterfactual inference can be performed upon such a model to sample harmonized data by intervening upon these variables. We demonstrated experimental results on two tasks, age regression and Alzheimer's Disease classification, on a wide range of real-world datasets. We showed that our method compares favorably to state-of-the-art algorithms such as IRM and ComBat. Future directions for this work include causal graph identification and causal mediation.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgements.

We thank Ben Glocker, Nick Pawlowski and Daniel C. Castro for suggestions. This work was supported by the National Institute on Aging (grant numbers RF1AG054409 and U01AG068057) and the National Institute of Mental Health (grant number R01MH112070). Pratik Chaudhari would like to acknowledge the support of the Amazon Web Services Machine Learning Research Award.

References

1. Arjovsky M, Bottou L, Gulrajani I, Lopez-Paz D: Invariant risk minimization. arXiv preprint arXiv:1907.02893 (2019)
2. Armstrong NM, An Y, Beason-Held L, Doshi J, Erus G, Ferrucci L, Davatzikos C, Resnick SM: Predictors of neurodegeneration differ between cognitively normal and subsequently impaired older adults. *Neurobiol. Aging* 75, 178–186 (2019) [PubMed: 30580127]
3. Bashyam VM, et al. : Medical image harmonization using deep learning based canonical mapping: Toward robust and generalizable learning in imaging. arXiv preprint arXiv:2010.05355 (2020)
4. Bingham E, et al. : Pyro: deep universal probabilistic programming. *J. Mach. Learn. Res* 20(1), 973–978 (2019)
5. Chen AA, Beer JC, Tustison NJ, Cook PA, Shinohara RT, Shou H: Removal of scanner effects in covariance improves multivariate pattern analysis in neuroimaging data. *bioRxiv*, p. 858415 (2020)
6. Davatzikos C: Machine learning in neuroimaging: progress and challenges. *Neuroimage* 197, 652 (2019) [PubMed: 30296563]
7. Dolatabadi HM, Erfani S, Leckie C: Invertible generative modeling using linear rational splines. arXiv preprint arXiv:2001.05168 (2020)
8. Doshi J, Erus G, Ou Y, Gaonkar B, Davatzikos C: Multi-atlas skull-stripping. *Acad. Radiol* 20(12), 1566–1576 (2013) [PubMed: 24200484]
9. Doshi J, et al. : Muse: multi-atlas region segmentation utilizing ensembles of registration algorithms and parameters, and locally optimal atlas selection. *Neuroimage* 127, 186–195 (2016) [PubMed: 26679328]
10. Durkan C, Bekasov A, Murray I, Papamakarios G: Neural spline flows. In: *Advances in Neural Information Processing Systems*, pp. 7511–7522 (2019)
11. Esteva A, et al. : Dermatologist-level classification of skin cancer with deep neural networks. *Nature* 542(7639), 115–118 (2017) [PubMed: 28117445]
12. Goodfellow I, et al. : Generative adversarial nets. In: *Advances in Neural Information Processing Systems*, pp. 2672–2680 (2014)
13. Habes M, et al. : The brain chart of aging: machine-learning analytics reveals links between brain aging, white matter disease, amyloid burden, and cognition in the iSTAGING consortium of 10,216 harmonized MR scans. *Alzheimer's Dement.* 17(1), 89–102 (2021) [PubMed: 32920988]
14. Hegenscheid K, Kühn JP, Völzke H, Biffar R, Hosten N, Puls R: Whole-body magnetic resonance imaging of healthy volunteers: pilot study results from the population-based ship study. In: *RöFo-Fortschritte auf dem Gebiet der Röntgenstrahlen und der bildgebenden Verfahren*, vol. 181, pp. 748–759. Georg Thieme Verlag KG Stuttgart · New York (2009)
15. Jack CR Jr., et al. : The Alzheimer's disease neuroimaging initiative (adni): mri methods. *J. Magn. Reson. Imaging Off. J. Int. Soc. Magn. Reson. Med* 27(4), 685–691 (2008)
16. Johnson WE, Li C, Rabinovic A: Adjusting batch effects in microarray expression data using empirical bayes methods. *Biostatistics* 8(1), 118–127 (2007) [PubMed: 16632515]
17. Kingma DP, Ba J: Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980 (2014)
18. Kingma DP, Welling M: Auto-encoding variational bayes. arXiv preprint arXiv:1312.6114 (2013)
19. Menze BH, et al. : The multimodal brain tumor image segmentation benchmark (brats). *IEEE Trans. Med. Imaging* 34(10), 1993–2024 (2014) [PubMed: 25494501]
20. Morris CN: Parametric empirical Bayes inference: theory and applications. *J. Am. Stat. Assoc* 78(381), 47–55 (1983)
21. Moyer D, Gao S, Brekelmans R, Galstyan A, Ver Steeg G: Invariant representations without adversarial training. *Adv. Neural. Inf. Process. Syst* 31, 9084–9093 (2018)
22. Moyer D, Ver Steeg G, Tax CM, Thompson PM: Scanner invariant representations for diffusion MRI harmonization. *Magn. Reson. Med* 84(4), 2174–2189 (2020) [PubMed: 32250475]
23. Papamakarios G, Nalisnick E, Rezende DJ, Mohamed S, Lakshminarayanan B: Normalizing flows for probabilistic modeling and inference. arXiv preprint arXiv:1912.02762 (2019)

24. Papamakarios G, Pavlakou T, Murray I: Masked autoregressive flow for density estimation. In: Advances in Neural Information Processing Systems, pp. 2338–2347 (2017)
25. Paszke A, et al. : Pytorch: an imperative style, high-performance deep learning library. Adv. Neural. Inf. Process. Syst 32, 8026–8037 (2019)
26. Pawlowski N, Castro DC, Glocker B: Deep structural causal models for tractable counterfactual inference. arXiv preprint arXiv:2006.06485 (2020)
27. Pearl J: Causality: Models, Reasoning, and Inference, 2nd edn. Cambridge University Press, Cambridge (2009)
28. Pearl J, et al. : Causal inference in statistics: an overview. Stat. Surv 3, 96–146 (2009)
29. Peters J, Janzing D, Schölkopf B: Elements of Causal Inference. The MIT Press, Cambridge (2017)
30. Pomponio R, et al. : Harmonization of large mri datasets for the analysis of brain imaging patterns throughout the lifespan. NeuroImage 208, 116450 (2020) [PubMed: 31821869]
31. Resnick SM, Pham DL, Kraut MA, Zonderman AB, Davatzikos C: Longitudinal magnetic resonance imaging studies of older adults: a shrinking brain. J. Neurosci 23(8), 3295–3301 (2003) [PubMed: 12716936]
32. Robinson R, et al.: Image-level harmonization of multi-site data using image-and-spatial transformer networks. In: Martel AL, et al. (eds.) MICCAI 2020. LNCS, vol. 12267, pp. 710–719. Springer, Cham (2020). 10.1007/978-303-0-59728-369
33. Schölkopf B: Causality for machine learning. arXiv preprint arXiv:1911.10500 (2019)
34. Sudlow C, et al. : UK biobank: an open access resource for identifying the causes of a wide range of complex diseases of middle and old age. Plos Med. 12(3), e1001779 (2015) [PubMed: 25826379]
35. Tustison NJ, et al. : N4itk: improved n3 bias correction. IEEE Trans. Med. Imaging 29(6), 1310–1320 (2010) [PubMed: 20378467]
36. Wachinger C, Rieckmann A, Pölsterl S, Initiative, A.D.N., et al. : Detect and correct bias in multi-site neuroimaging datasets. Med. Image Anal 67, 101879 (2021) [PubMed: 33152602]
37. Zhu JY, Park T, Isola P, Efros AA: Unpaired image-to-image translation using cycle-consistent adversarial networks. In: Proceedings of the IEEE International Conference on Computer Vision, pp. 2223–2232 (2017)




Fig. 1.

Causal graph of the structural causal model for brain imaging. Data consists of brain ROIs (x), sex (s), age (a), imaging site (t), and their respective exogenous variables (ϵ_x , ϵ_s , ϵ_a , and ϵ_t). Bidirectional arrows indicate invertible normalizing flow models and the black dot shows that the flow model associated with x is conditioned on the direct causes (parents) s , a , and t ; this follows the notation introduced in [26]. We are interested in answering counterfactual questions of the form “what would the scans look like if they had been acquired from the same site”. We first train the flow-based SCM M_θ on the observed data. We then infer the posterior of exogenous variables ϵ_x and ϵ_a with the invertible structural assignments (abduction step). We can now intervene upon the site by replacing site variable

t with a specific value τ , this is denoted by $\text{do}(t = \tau)$. We sample from the modified flow-based SCM $M_{\text{do}(t=\tau)}$ to obtain counterfactual queries.




Fig. 2. Comparison of normalized feature (hippocampus volume) distributions for various sites in the iSTAGING consortium data before (raw) and after harmonization using ComBat-Linear and our SCM-QSpline. We observe that ComBat aligns inter-site feature distributions by preserving sex and age effects and removes all other unknown confounders by treating them as site effects. In contrast, the distribution of hippocampus volume is unchanged in our proposed method which takes both known confounders (sex, age, and site) and unknown confounders (exogenous noises) into consideration for harmonization. ComBat removing these useful confounders is detrimental to accuracy (see Table 2).

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

Table 1.

Mean average error of age prediction for data from the iSTAGING consortium.

	Study	TarOnly	SrcOnly	IRM	ComBat				Flow-based SCM (ours)		
					Linear	GAM	ComBat+ +	CovBat	Affine	L-Spline	Q-Spline
Source	BLSA-3T	–	11.74 (0.35)	11.76 (0.35)	11.72 (0.62)	11.74 (0.61)	11.73 (0.62)	11.74 (0.62)	11.74 (0.61)	11.74	11.65 (0.62)
Target	BLSA-1.5T	6.77 (0.82)	7.21 (0.91)	7.16 (0.87)	7.14 (0.99)	7.01 (0.99)	7.00 (1.04)	7.03 (1.08)	7.01 (1.01)	7.00 (1.04)	6.92 (1.09)
Target	UKBB	6.14 (0.16)	7.27 (0.70)	7.18 (0.58)	6.62 (0.46)	6.70 (0.46)	6.71 (0.47)	6.75 (0.49)	6.72 (0.46)	6.75 (0.47)	6.44 (0.28)
Target	SHIP	11.36 (0.31)	17.14 (0.62)	17.05 (0.46)	15.95 (0.61)	16.17 (0.59)	16.21 (0.47)	16.22 (0.65)	16.20	16.25 (0.63)	15.68 (0.80)

All experiments were repeated 5 times in cross-validation fashion, and the average performance is reported with the standard errors in the brackets. TarOnly indicates validation MAEs directly trained on each target sites. The hypothesis that our proposed methods achieve a better accuracy than the baselines can be accepted with p -values between 0.06–0.41. This task can be interpreted as a sanity check for our method.

Table 2.

AD classification accuracy (%) comparison on the ADNI dataset and standard deviation (in brackets) across 5-fold cross-validation.

	Study	TarOnly	SrcOnly	IRM	ComBat				Flow-based SCM (ours)		
					Linear	GAM	ComBat+ +	CovBat	Affine	L-Spline	Q-Spline
Source	ADNI-1	–	76.1 (1.54)	76.2 (2.46)	75.1 (1.37)	75.1 (1.23)	65.1 (6.29)	74.4 (2.29)	76.1 (1.92)	75.3 (1.76)	75.4 (2.45)
Target	ADNI-2	75.8 (3.46)	71.9 (4.88)	73.0 (4.85)	71.4 (4.30)	72.1 (2.83)	56.2 (9.29)	67.4 (5.06)	73.4 (3.52)	72.6 (3.48)	73.7 (4.13)
Source	ADNI-2	–	75.8 (3.46)	76.3 (2.35)	77.5 (2.30)	77.0 (2.74)	67.8 (9.42)	77.9 (2.47)	78.7(1.32)	78.2 (2.80)	77.5(1.76)
Target	ADNI-1	76.1 (1.54)	70.4 (8.80)	72.0 (2.16)	71.1 (4.07)	70.1 (5.67)	58.0 (6.28)	69.1 (5.82)	71.4 (2.41)	71.8 (5.76)	73.3 (3.04)

TarOnly indicates validation accuracies for training directly on each target site. The hypothesis that our proposed method (Q-Spline) achieves a better accuracy than baselines can be accepted with p -values less than 10^{-5} .