

Detect and Correct Bias in Multi-Site Neuroimaging Datasets

Christian Wachinger^{1*}, Anna Rieckmann², Sebastian Pölsterl¹
for the Alzheimer's Disease Neuroimaging Initiative
and the Australian Imaging Biomarkers and Lifestyle flagship study of ageing

¹Lab for Artificial Intelligence in Medical Imaging (AI-Med), Department of Child and Adolescent Psychiatry, University Hospital, LMU München, Germany

²Umeå Center for Functional Brain Imaging, Department of Radiation Sciences, Umeå University, Sweden

Abstract

The desire to train complex machine learning algorithms and to increase the statistical power in association studies drives neuroimaging research to use ever-larger datasets. The most obvious way to increase sample size is by pooling scans from independent studies. However, simple pooling is often ill-advised as selection, measurement, and confounding biases may creep in and yield spurious correlations. In this work, we combine 35,320 magnetic resonance images of the brain from 17 studies to examine bias in neuroimaging. In the first experiment, *Name That Dataset*, we provide empirical evidence for the presence of bias by showing that scans can be correctly assigned to their respective dataset with 71.5% accuracy. Given such evidence, we take a closer look at confounding bias, which is often viewed as the main shortcoming in observational studies. In practice, we neither know all potential confounders nor do we have data on them. Hence, we model confounders as unknown, latent variables. Kolmogorov complexity is then used to decide whether the confounded or the causal model provides the simplest factorization of the graphical model. Finally, we present methods for dataset harmonization and study their ability to remove bias in imaging features. In particular, we propose an extension of the recently introduced ComBat algorithm to control for global variation across image features, inspired by adjusting for unknown population stratification in genetics. Our results demonstrate that harmonization can reduce dataset-specific information in image features. Further, confounding bias can be reduced and even turned into a causal relationship. However, harmonization also requires caution as it can easily remove relevant subject-specific information. Code is available at <https://github.com/ai-med/Dataset-Bias>.

Keywords: Bias, MRI, Big Data, Causal inference, Harmonization

1. Introduction

Is it possible to predict the dataset that a brain scan is coming from based on image-derived measures like volume or thickness? Initially, we may guess that it should be impossible. On second thought, we may notice that it depends on the datasets we are comparing. If one dataset

only contains adolescent subjects and another one only elderly subjects, it should be possible to distinguish them due to the association of image-derived measures with age. But next to demographics, are there unique signatures in each dataset that would facilitate identifying the source of an image? We will demonstrate that the source dataset can indeed be identified with much higher accuracy than would be expected from basic demographics of the participants. Hence, not only subject-specific, but also imaging site-specific information is implicitly encoded. This is insofar surprising as we are working with image-derived measures, which should only relate to the underlying neuroanatomy of the subject and not to the imaging site, where the scan was acquired.

Colloquially, this phenomenon is referred to as *dataset bias*. In statistics, bias refers to a systematic deviation from the true, possibly unknown, underlying quantitative parameter that is being estimated. In a typical neuroimaging study, various types of bias can be present that can alter the conclusions one deduces from this study. In the first step, individuals have to be enrolled into the study. If subjects do not faithfully represent the overall popula-

*Corresponding Author. Address: Waltherstr. 23, 80337 München, Germany; Email: christian.wachinger@med.uni-muenchen.de

**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/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 (www.loni.usc.edu/ADNI). The AIBL researchers contributed data but did not participate in analysis or writing of this report. AIBL researchers are listed at www.aibl.csiro.au.

tion one wants to study, i.e., one obtains a non-random sample of a population, conclusions will be biased. This is referred to as *selection bias*. For instance, selection bias is present if the study recruits particular target groups, e.g., young adults or patients with a particular disease. While such a selection may be related to the study objective and therefore done on purpose, there are also hidden, unintended factors like the over-representation of more educated participants (Smith and Nichols, 2018) that would yield biased estimates with respect to the overall population.

After enrollment, subjects will undergo magnetic resonance imaging. Prior studies on inter-scanner variability have already noted that there is a strong dependence of the acquired image on magnetic field strength, manufacturer, gradients, pulse sequences, and head positioning (Jovicich et al., 2009). While standardization efforts are undertaken for instance by the ADNI (Jack et al., 2008), variations related to the scanner remain (Krugger et al., 2010), and it is even questionable if a further standardization is in the manufacturer’s interest. Even when assuming a faithful image reconstruction, scans often undergo various image analysis steps to derive summary statistics that depend on the algorithm being used. The choice of segmentation and registration algorithms can cause varying results, also subject to the input, potentially affected by motion artifacts, voxel sizes, and image noise. Therefore, neuroimaging data is unavoidably subject to various types of *measurement biases*.

Once data has been collected and processed, usually image-derived measures are related to disease status or outcomes of a neurocognitive test to determine which brain structures are responsible for the observed outcome. It is important to remember that, in general, regression analysis can only establish correlation, but not causation. Association, unlike causation, is a symmetric relationship: two variables X and Y are associated, regardless of whether X causes Y , or Y causes X . Only under exchangeability, correlation implies causation (Hernan and Robins, 2020). A common violation of exchangeability is the presence of confounding. For example, consider you want to study causes of Alzheimer’s disease (AD). Analyses of the data show a high correlation between gray hair and AD, which may naïvely lead to the conclusion that gray hair causes AD. However, the observed correlation between gray hair and AD is only due to a person’s age. Therefore, the association between gray hair and AD is confounded by the common cause age. This form of bias is known as *confounding bias*.

Dataset bias is becoming pivotal as neuroimaging is joining the ranks of a “big data” science with more and larger datasets becoming available (Smith and Nichols, 2018). Large sample sizes are required for a number of applications in neuroimaging, such as association studies in imaging-genetics, or training complex machine learning models – in particular in deep learning. As outlined above, sources of bias are plentiful and depend on the re-

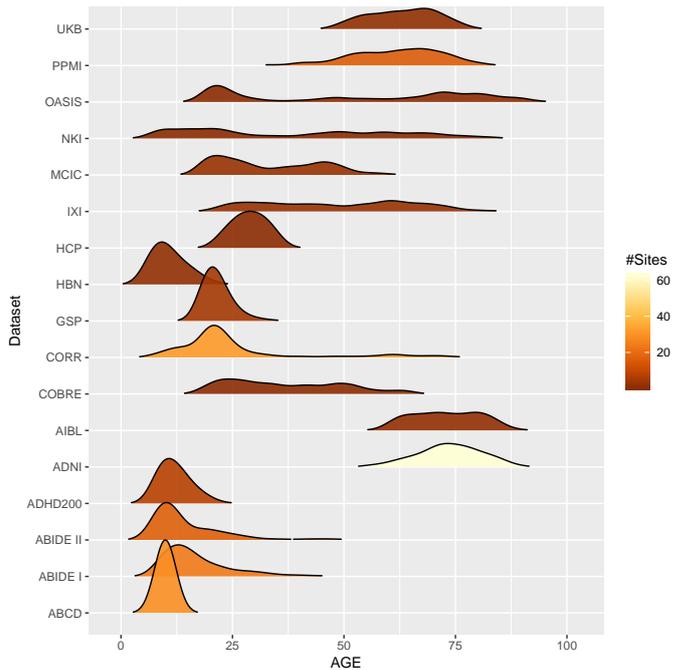


Figure 1: Age distribution per dataset. The fill color corresponds to the number of imaging sites per dataset.

search question and the data that is used to answer it. For this reason, neuroimaging data is usually collected with a particular research question in mind, and inclusion criteria are tailored to answering this particular question as unbiased as possible, e.g., by randomization or collecting information on possible confounders. Therefore, pooling data from studies that have been designed with different research questions in mind, will likely lead to bias in a machine learning model trained on this data. In contrast, if a model would be truly unbiased on a population level, it would naturally generalize to other datasets.

In this article, we study bias in neuroimaging data. To this end, we combine data from 17 large-scale studies, presented in section 2. First, we propose *Name That Dataset*, i.e., the prediction of the dataset that a scan is part of, as an experiment to detect inter-dataset bias in section 3. Second, we take a closer look at confounding bias in section 4 and present a method for distinguishing between causal and confounded statistical relationships in a single dataset using causal inference. As we do not know all potential confounders, we work with a model that assumes an unobserved confounding variable. Third, we discuss methods for dataset harmonization as a means of reducing bias in the data, and present an extension that accounts for global variation across features in section 5. Finally, we present results for dataset prediction on harmonized data, brain age prediction with harmonization, and the impact of harmonization on two causal models in section 6. An earlier version of this work has been presented at a conference (Wachinger et al., 2019) and is extended in this

Table 1: Overview of neuroimaging datasets used in this study.

Dataset	Diagnosis	N	Age (mean)	Age (SD)	Males %	Sites	Diseased
ABCD		8,751	9.9	0.6	51.6	29	-
ABCD Subset		(1,570)	9.9	0.6	51.6	29	-
ABIDE I	Autism	1,095	17.1	8.1	85.2	24	526
ABIDE II	Autism	1,032	15.2	9.4	76.1	17	477
ADHD200	ADHD	965	12.1	3.3	61.8	8	384
ADNI	Alzheimer’s	1,682	73.6	7.2	55.0	62	1,144
AIBL	Alzheimer’s	262	72.9	7.6	47.3	2	91
COBRE	Schizophrenia	146	37.0	12.8	74.7	1	72
CORR		1,476	25.9	15.4	50.0	32	0
GSP		1,563	21.5	2.8	42.3	5	0
HBN	Psychiatric	689	10.7	3.6	59.7	2	497
HCP		1,113	28.8	3.7	45.6	1	0
IXI		561	48.6	16.5	44.6	3	0
MCIC	Schizophrenia	194	33.1	11.5	71.6	3	104
NKI	Psychiatric	624	38.4	22.5	39.1	1	268
OASIS	Alzheimer’s	415	52.8	25.1	38.6	1	100
PPMI	Parkinson’s	390	61.2	10.0	62.6	16	284
UKB	Diverse	14,362	62.8	7.5	47.6	2	1,878
UKB Subset	Diverse	(1,474)	63.2	7.4	47.4	2	160
		35,320	39.6	25.9	51.6	210	

article with more datasets, details, experiments, and the addition of harmonization.

2. Neuroimaging Data

We work on MRI T1 brain scans from 17 large-scale public datasets:

1. Adolescent Brain Cognitive Development (ABCD) (Casey et al., 2018),
2. Autism Brain Imaging Data Exchange (ABIDE) I (Di Martino et al., 2014),
3. Autism Brain Imaging Data Exchange (ABIDE) II (Di Martino et al., 2017),
4. Attention Deficit Hyperactivity Disorder (ADHD200) (Milham et al., 2012),
5. Alzheimer’s Disease Neuroimaging Initiative (ADNI) (Jack et al., 2008) (for up-to-date information, see www.adni-info.org),
6. Australian Imaging Biomarkers and Lifestyle Study of Ageing (AIBL) (Ellis et al., 2009),
7. Center for Biomedical Research Excellence (COBRE) (Mayer et al., 2013),
8. Consortium for Reliability and Reproducibility (CORR) (Zuo et al., 2014),
9. Genomic Superstruct Project (GSP) (Buckner et al., 2012),
10. Healthy Brain Network (HBN) (Alexander et al., 2017),
11. Human Connectome Project (HCP) (Van Essen et al., 2013),
12. IXI Dataset,¹

13. Open Access Series of Imaging Studies (OASIS) cross-sectional sample (Marcus et al., 2007),
14. MIND Clinical Imaging Consortium (MCIC) schizophrenia sample (Gollub et al., 2013),
15. Nathan Kline Institute – Rockland Sample (NKI) (Nooner et al., 2012),
16. Parkinson Progression Marker Initiative (PPMI) (Marek et al., 2011),
17. UK Biobank Imaging (UKB) (Miller et al., 2016).

All datasets were processed with FreeSurfer (Fischl et al., 2002) version 5.3. We keep only one scan per subject from longitudinal or test-retest datasets. After exclusion of scans with processing errors and incomplete meta-data, scans from 35,320 subjects remained. Fig. 1 illustrates the age distribution per dataset together with the number of sites. Table 1 presents an overview of demographics per dataset. The ABCD and UKB datasets stand out as the largest datasets, which introduces a severe class imbalance for the dataset prediction experiment in section 3. Hence, we downsampled ABCD and UKB data to obtain a sample size close to the next largest datasets. We preserved the heterogeneity of the datasets by ensuring that all imaging sites are included in the subset. For UKB, we marked subjects with psychiatric diseases and cancer in the head as diseased. For ABCD, the diagnostic information is highly complex as mental health of adolescents is investigated in general, so that we have not encoded it for this study.

3. Name That Dataset

In order to detect the impact of dataset bias, we play the game *Name That Dataset* on neuroimaging data that was

¹<http://brain-development.org/ixi-dataset/>

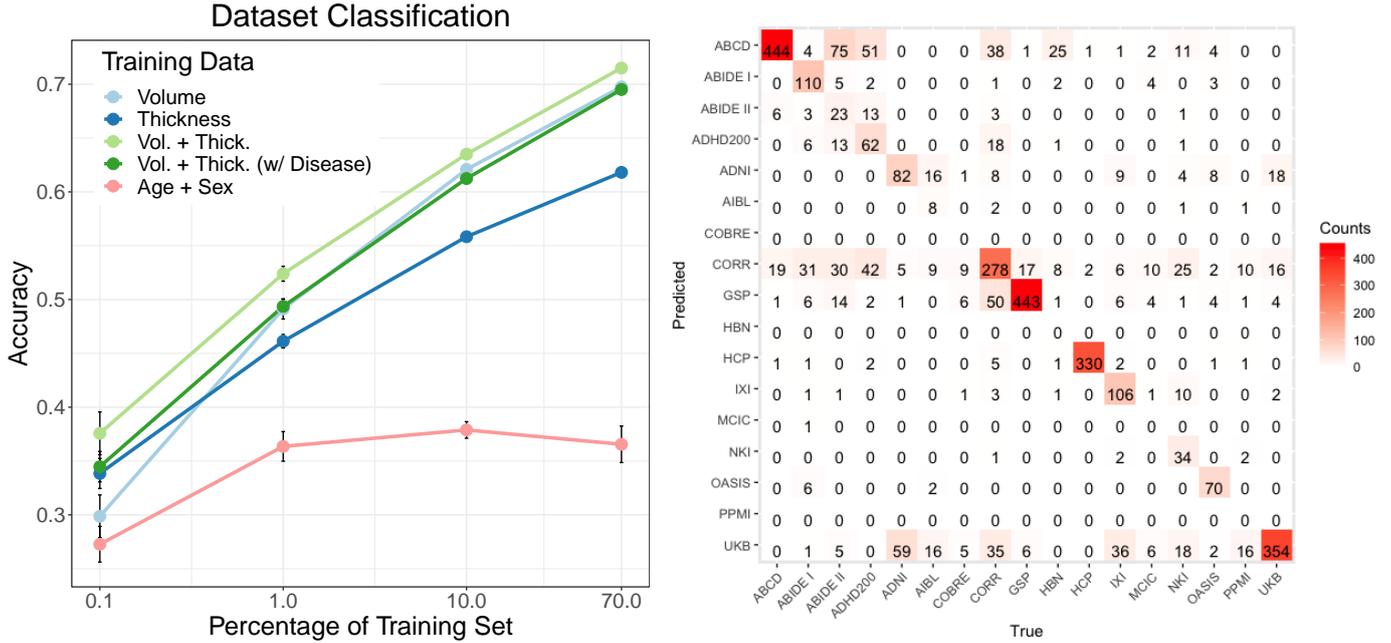


Figure 2: Left: Dataset classification accuracy for age and sex, volume, thickness, and their combination. The percentage of the data used for training is shown in log-scale. Lines show the average score over 50 repetitions, error bars show the standard deviation. Right: Confusion matrix for volume and thickness with 70% training data.

originally proposed by Torralba and Efros (2011) on natural images. The task is to predict the dataset D an MRI scan is coming from solely based on image-derived measures Ψ , where we use volume and thickness.² A random forest classifier is trained for the prediction $\Psi \mapsto D$, where we use default settings from the R package by Liaw et al. (2002).³ If dataset bias would be absent, we would expect a prediction accuracy close to random chance, which would be 5.9% for 17 datasets; taking unequal dataset sizes into account, it would be 7.8%. As additional reference, we report the results of a classifier trained on age and sex alone, which can achieve higher accuracy than random chance as datasets focus on specific parts of the population. In the experiments, we split data into training and testing sets with stratified sampling by dataset. This ensures that also the small datasets are represented in the training set for low sampling rates.

Fig. 2 illustrates the performance for classifying images into one of 17 datasets from volume and thickness measures, and their combination. We mainly report results on healthy controls to exclude disease-specific effects that could potentially ease classification; for the combination of volume and thickness we also show results for including diseased subjects, indicated as ‘(w/ Disease)’. In this experiment, we vary the amount of training data from 0.1% to 70%. The classification accuracy widely improves as the

training set increases, particularly for the image features. The highest accuracy of 71.5% is achieved by the combination of volume and thickness measures, which perform better than each of them alone. Except for the 0.1% training set size, volumes have a higher accuracy than thickness measures; for 70% training, the difference in accuracy is 7.9%. The accuracy after the inclusion of diseased patients is slightly lower compared to healthy subjects, as shown for the combination of volume and thickness. The classifier with age and sex reaches 37.9% accuracy, which is well above random and therefore hints at selection bias. Yet, compared to over 70% accuracy for image features, there must be another source of bias that cannot be explained by basic demographics, such as measurement and confounding bias.

From the confusion matrix in Fig. 2, the high classification accuracy (diagonal elements) indicates that datasets possess unique, identifiable characteristics. In addition, we can see that datasets with a similar population result in higher confusion, e.g., between ABCD, ABIDE I+II, and ADHD200. Single-site datasets, like HCP, that have strict inclusion criteria and imaging protocols show almost no confusion with any of the other datasets. In contrast, multi-site datasets like CORR that also cover a wide age range, show high confusion with other datasets. Also for UKB, with its large size and age distribution, we observe confusion with other datasets, although scans were only acquired at two sites with the same scanner.

The lesson learned from this experiment is that even when working with image-derived values that represent

²We used 55 volume measures from FreeSurfer as listed in aseg.stats and 70 mean thickness measures of cortical parcels as listed in lh.aparc.stats and rh.aparc.stats.

³500 trees, mtry is square root of number of variables

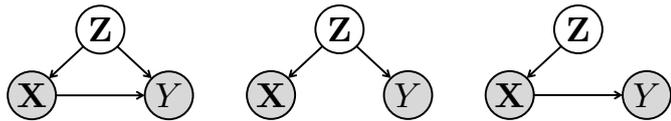


Figure 3: Probabilistic graphical models for observed variables \mathbf{X} , Y and unobserved confounders \mathbf{Z} . The statistical relationship between \mathbf{X} and Y is due to confounder \mathbf{Z} and due to the influence of \mathbf{X} on Y (left). Limiting cases are pure confounding (middle) and pure causality (right).

physical measures (volume, thickness), substantial bias in datasets remains, despite techniques like atlas renormalization (Han and Fischl, 2007) were employed to improve consistency across scanners. Of course, much of the bias can be attributed to the different goals of the studies, like the inclusion of subjects from different age groups. However, even when focusing on datasets that cover a similar age range, we observe a high accuracy, e.g., ABIDE I and II. While we are not aware of previous attempts to *Name That Dataset*, our results echo concerns raised in previous studies. In an ENIGMA study with over 15,000 subjects on brain asymmetry (Guadalupe et al., 2017), it was reported that dataset heterogeneity explained over 10% of the total observed variance per structure. In ADNI, which has an optimized MPRAGE imaging protocol across all sites (Jack et al., 2008), the intra-subject variability of compartment volumes for scans on different scanners was roughly 10 times higher than repeated scans on the same scanner (Krugger et al., 2010). Similarly, Wachinger and Reuter (2016) reported a drop in accuracy for Alzheimer’s disease prediction when training and testing on different datasets.

4. Telling Causal from Confounded with Causal Inference

In the previous section, we have established that there is correlation between a feature vector ψ , derived from MRI scans, and the dataset D the scan belongs to, by estimating the probability $P(D = d | \Psi = \psi)$ via a random forest classifier. While this has yielded useful insights, it only provides a measure of statistical dependence, which alone is insufficient to determine causal structures of confounding bias.

4.1. Confounding Bias in Causal Inference

Here, we want to study bias in a more principled manner by looking at confounding bias in a causal inference framework. Generally, it is challenging to infer cause and effect from observational data, as this normally requires randomized controlled trials (Pearl, 2000). The problem is that it is difficult to determine whether a variable X causes Y or whether both variables have a common cause Z . If X and Y share the cause Z , then confounding is present and Z is referred to as confounder, as illustrated in Fig. 3. Because there is an open backdoor path between

X and Y , it is unclear if an association between X and Y results from the causal effect $X \rightarrow Y$ or from the path with the common cause $X \leftarrow Z \rightarrow Y$. The extreme cases would be the purely causal setting without confounding $X \rightarrow Y$, illustrated in Fig. 3 (right) and the purely confounded setting, where the correlation between X and Y is entirely due to common cause Z , illustrated in Fig. 3 (middle).

For inferring causality, one of the core assumptions is *causal sufficiency*, i.e., we know all potential confounders. This assumption is often violated in practice, since we do not know all potentially relevant factors nor do we have data on them, yielding spurious inferences. Even if they would be known, they would not be available in all publicly available neuroimaging datasets. Consequently, we consider a recently proposed approach for causal inference by Kaltenpoth and Vreeken (2019) that assumes an *unobserved* confounding variable. The method explicitly models the hidden confounder with probabilistic PCA, which allows comparing the causal $X \rightarrow Y$ and confounded model $X \leftarrow Z \rightarrow Y$ to conclude whether the relationship is causal or confounded. Next, we will present the details of the approach.

4.2. Assumptions for Inferring Causality with Unknown Confounders

We consider samples from the joint distribution $P(\mathbf{X}, Y)$ over two statistically dependent continuous-valued random variables, where \mathbf{X} can be multivariate and Y is univariate. The aim is to determine whether it is more likely that \mathbf{X} causes Y , or that there exists an unobserved random variable \mathbf{Z} that is the cause of both \mathbf{X} and Y . \mathbf{Z} can be multivariate. Inferring causal relations from observational data is only attainable under specific model assumptions. A general assumption in causal inference is that the data distribution can be represented as a causal directed acyclic graph that satisfies the Markov factorization property: conditioned on its parents (direct causes), each variable is independent of its nondescendants (Pearl, 2000). In addition, faithfulness and the previously mentioned assumption of causal sufficiency are required. Faithfulness implies that if \mathbf{X} is independent of Y , there is no direct influence between the two in the underlying graph (Pearl, 2000). While faithfulness is a strong condition, it is generally reasonable. Establishing causal sufficiency is more problematic in practice.

Our approach incorporates the algorithmic Markov condition (AMC) (Janzing and Scholkopf, 2010), which states that if \mathbf{X} causes Y , the factorization of the joint distribution $P(\mathbf{X}, Y)$ in the true causal direction has a lower Kolmogorov complexity than in the anti-causal direction. Together with causal sufficiency, this allows to identify the true causal network as the least complex one. Although the AMC generally relies on causal sufficiency, Kaltenpoth and Vreeken (2019) proposed to incorporate the confounder as a latent variable \mathbf{Z} . This approach allows for computing whether confounding is present, with-

out explicitly knowing the confounder or having data on it, which is very relevant in practice.

A remaining challenge for estimating the AMC is the evaluation of the Kolmogorov complexity as it is not directly computable. To this end, the minimum description length (MDL) principle provides a statistically well-founded approach to approximate Kolmogorov complexity (Grünwald, 2007). Considering model class \mathcal{M} and a fully Bayesian formulation, the code length function L is computed as

$$L(\mathbf{X}|\mathcal{M}) = -\log \int_{M \in \mathcal{M}} P(\mathbf{X}|M) dP(M), \quad (1)$$

where $P(M)$ is a prior on the model class \mathcal{M} .

4.3. Causal and Confounded Models

Now that we established the theoretical foundation of our work, we need to define two factorizations of $P(\mathbf{X}, Y)$ – one for the causal and one for the confounded scenario – to decide whether the relationship between \mathbf{X} and Y is indeed causal or not. Figure 3 illustrates the two factorizations we consider here:

1. Causal: $P(\mathbf{X}, Y) = P(Y|\mathbf{X})P(\mathbf{X}|\mathbf{Z})P(\mathbf{Z})$, which is represented by a linear regression model,
2. Confounded: $P(\mathbf{X}, Y) = P(Y|\mathbf{Z})P(\mathbf{X}|\mathbf{Z})P(\mathbf{Z})$, which is estimated by probabilistic PCA (Tipping and Bishop, 1999).

The complexity under the causal model can be estimated by minimum description length $L_{ca}(\mathbf{X}, Y)$ via

$$\begin{aligned} L_{ca}(\mathbf{X}, Y) &= -\log P(\mathbf{X}) \int P(Y|\mathbf{X}, \mathbf{w})P(\mathbf{w})d\mathbf{w}, \\ \mathbf{X} &\sim \mathcal{N}(\mathbf{0}, \sigma_x^2 I), \\ \mathbf{w} &\sim \mathcal{N}(\mathbf{0}, \sigma_w^2 I), \\ Y|\mathbf{X}, \mathbf{w} &\sim \mathcal{N}(\mathbf{w}^\top \mathbf{X}, \sigma_y^2 I). \end{aligned} \quad (2)$$

The complexity of the confounded model is estimated by

$$\begin{aligned} L_{co}(\mathbf{X}, Y) &= \\ &-\log \int P(\mathbf{X}, Y|\mathbf{Z}, \mathbf{W})P(\mathbf{Z})P(\mathbf{W})d\mathbf{W}d\mathbf{Z}, \\ \mathbf{Z} &\sim \mathcal{N}(\mathbf{0}, \sigma_z^2 I), \\ [\mathbf{W}, \mathbf{w}_y] &\sim \mathcal{N}(\mathbf{0}, \sigma_w^2 I), \\ \mathbf{X}|\mathbf{Z}, \mathbf{W} &\sim \mathcal{N}(\mathbf{W}^\top \mathbf{Z}, \sigma_x^2 I), \\ Y|\mathbf{Z}, \mathbf{W} &\sim \mathcal{N}(\mathbf{w}_y^\top \mathbf{Z}, \sigma_y^2 I), \end{aligned} \quad (3)$$

where the confounders \mathbf{Z} and the principal axes \mathbf{W} , \mathbf{w}_y are estimated using probabilistic PCA by appending Y as an extra column to \mathbf{X} (Tipping and Bishop, 1999).

Note that we do not require that the confounders are known or measured; since we marginalize over \mathbf{Z} , we only need to specify its dimensionality k . To compare the causal

($\mathbf{X} \rightarrow Y$) and the confounded model ($\mathbf{X} \leftarrow \mathbf{Z} \rightarrow Y$), we compute

$$\Delta(\mathbf{X}, Y) = L_{co}(\mathbf{X}, Y) - L_{ca}(\mathbf{X}, Y). \quad (4)$$

If the causal model better describes the data than the confounded model, we obtain $\Delta(\mathbf{X}, Y) > 0$ – the more positive, the more confident we are. If instead $\Delta(\mathbf{X}, Y) < 0$, the roles are reversed. As Δ is dependent on the size of the dataset, we also consider the normalized version $\frac{1}{N}\Delta$, with N denoting the number of subjects.

We first estimate the posterior distribution of the model parameters $\sigma_x, \sigma_y, \mathbf{w}_y, \mathbf{W}, \mathbf{Z}$ via Markov chain Monte Carlo using a no-U-turn sampler (Hoffman and Gelman, 2014) with four chains and 8,000 samples each (of which 1,500 are used for burn-in) as implemented in RStan 2.19.3 (Carpenter et al., 2017; Stan Development Team, 2020). We experienced instability estimating the posterior using variational inference as proposed by Kaltentpoth and Vreeken (2019). The integral over the model parameters in the marginal likelihood in Eqs. (2) and (3) is analytically intractable and we need to approximate it via Monte Carlo sampling. Instead of using the naïve Monte Carlo estimator of the marginal likelihood, we estimate it via bridge sampling, which has lower variance (Frühwirth-Schnatter, 2004; Gronau et al., 2020). We repeated this process ten times to obtain an interval of possible values for L_{ca} and L_{co} .

5. Data Harmonization

In the previous sections, we described methods for identifying confounding bias. Here, we want to study whether methods for harmonizing neuroimaging datasets can effectively reduce bias, and therefore can enable to identify causal effects. Harmonization of multi-site data aims to remove unwanted variability associated with scanner and site when pooling data. We refer to such unwanted variation as *non-biological variation* to distinguish it from *biological variation* about the subject, which we want to retain. Current methods for harmonization focus on adjusting for known confounders, but as argued in section 4, neither do we know all confounding factors nor do we have data on them. We avoid this issue by introducing a novel harmonization approach that can be used in the presence of unknown confounders by computing substitute confounders.

Harmonization can be performed at multiple stages: by specifying a strict acquisition protocol, by normalization in image processing, or by adjusting image-derived features before data analysis. Here, we will focus on harmonization at the feature stage, as we cannot influence the acquisition anymore, and we are using FreeSurfer for the image analysis, which already has a sophisticated normalization pipeline.

We will discuss several existing methods for harmonization, and propose our novel extension to account for unknown confounders below. Table 2 summarizes the dif-

ferent harmonization models and the corresponding variable updates. First, we need to specify on which level we want to harmonize features. For multi-site datasets, the obvious choice is to harmonize scans with respect to imaging site. Based on your initial motivation for *Name That Dataset*, an alternative would be to harmonize with respect to image source (dataset). We will evaluate both in our experiments, but focus on imaging site in the following description.

5.1. Harmonization only based on imaging sites

A naïve approach for harmonizing across imaging sites is **Z-score normalization**. Consider an image-derived measurement Y_{ijf} for imaging site i , subject j , and feature type f . After computing the mean $\hat{\mu}_{if}$ and standard deviation $\hat{\sigma}_{if}$ for each imaging site and feature, the feature’s harmonized value \mathcal{Y}_{ijf} is

$$\mathcal{Y}_{ijf} = \frac{Y_{ijf} - \hat{\mu}_{if}}{\hat{\sigma}_{if}}, \quad (5)$$

where subject j was scanned at site i . Z-score normalization is a standard procedure in statistics, simple to compute, but might remove information that is required for downstream tasks, such as disease prediction. Consider one imaging site with only young subjects and another one with only elderly subjects, both sites will have the same mean after the normalization, which for structures like the ventricles is inadequate (Scahill et al., 2003).

As second harmonization approach controls for site-specific effects by computing residuals. To this end, site becomes a regressor in a **linear regression** and the feature’s harmonized value is the residual:

$$Y_{ijf} = \alpha_f + \gamma_{if} + \varepsilon_{ijf}, \quad (6)$$

$$\mathcal{Y}_{ijf} = Y_{ijf} - \hat{\gamma}_{if}, \quad (7)$$

where α_f is the average measure for the reference site, γ_{if} is an additive imaging site effect, and ε_{ijf} is the variance. The parameters can be estimated by solving the corresponding ordinary least squares problem.

5.2. ComBat Harmonization

As an extension of the previous model, Fortin et al. (2018) proposed harmonization of cortical thickness measures based on Combining Batches (ComBat) (Johnson et al., 2007) from gene expression analysis. ComBat adds a site-specific scaling factor δ , yielding a model that adjusts for additive and multiplicative effects. In addition, ComBat uses empirical Bayes for inferring model parameters, which assumes that model parameters across features are drawn from the same distribution. Hence, model parameters are not estimated independently per feature anymore, which can be helpful with small sample sizes, but also assumes that parameters across features are homogeneous.

Accounting for imaging site can help in removing unwanted variation in the data, but at the same time it may

have the detrimental effect of removing relevant information. For instance, we would like to keep subject-specific information due to biological variability. We add the vector \mathbf{k} to the model to indicate variables, whose influence we would like to keep after harmonization. In our experiments, we use sex and a linear and quadratic age term to account for non-linear aging effects (Walhovd et al., 2011). The model is

$$Y_{ijf} = \alpha_f + \gamma_{if} + \beta_f^\top \mathbf{k}_j + \delta_{if} \varepsilon_{ijf} \quad (8)$$

with the multiplicative imaging site effect δ_{if} and vector of regression coefficients β_f . The harmonized value is computed as

$$\mathcal{Y}_{ijf} = \frac{Y_{ijf} - \hat{\alpha}_f - \hat{\beta}_f^\top \mathbf{k}_j - \hat{\gamma}_{if}}{\hat{\delta}_{if}} + \hat{\alpha}_f + \hat{\beta}_f^\top \mathbf{k}_j. \quad (9)$$

The linear model can also be adapted to accommodate variables to keep \mathbf{k} , as shown in Table 2.

5.3. Harmonization for non-biological variability with ComBat

As a novel extension, we propose to add the vector \mathbf{r} to the model that contains variables that relate to non-biological variability, which we would like to explicitly remove from the features. The updated ComBat model⁴ is

$$Y_{ijf} = \alpha_f + \gamma_{if} + \beta_f^\top \mathbf{k}_j + \zeta_f^\top \mathbf{r}_j + \delta_{if} \varepsilon_{ijf}, \quad (10)$$

with the vector of regression coefficients ζ_f . The harmonized values are

$$\mathcal{Y}_{ijf} = \frac{Y_{ijf} - \hat{\alpha}_f - \hat{\beta}_f^\top \mathbf{k}_j - \hat{\zeta}_f^\top \mathbf{r}_j - \hat{\gamma}_{if}}{\hat{\delta}_{if}} + \hat{\alpha}_f + \hat{\beta}_f^\top \mathbf{k}_j. \quad (11)$$

Known effects about non-biological variability that we would like to remove are, for instance, the scanner manufacturer (MF) or the magnetic field strength (MFS). However, we believe that there are also *unknown effects* that increase variability among sites, which we did not collect explicitly. Hence, we propose to adapt an approach that is common in genome-wide association studies (GWAS). A major concern in GWAS is population stratification, i.e., the existence of unobserved subpopulations in the sample. To account for population stratification in the analysis, the addition of principal components (PC) to the regression model has been proposed in EIGENSTRAT (Price et al., 2006). Principal components are computed on all genetic markers and therefore capture overall variation. Following this idea, we compute principal components *across all image features* on the whole dataset to capture generic variation that is not specific to a single brain feature. Recently, Wang and Blei (2019) proposed the deconfounder method

⁴The extension is available as ComBat++ in our repository <https://github.com/ai-med/Dataset-Bias>.

Variables		Model	Update
Site	Linear	$Y_{ijf} = \alpha_f + \gamma_{if} + \varepsilon_{ijf}$	$\mathcal{Y}_{ijf} = Y_{ijf} - \hat{\gamma}_{if}$
	ComBat	$Y_{ijf} = \alpha_f + \gamma_{if} + \delta_i \varepsilon_{ijf}$	$\mathcal{Y}_{ijf} = \frac{Y_{ijf} - \hat{\alpha}_f - \hat{\gamma}_{if}}{\hat{\delta}_i} + \hat{\alpha}_f$
Site & Keep	Linear	$Y_{ijf} = \alpha_f + \gamma_{if} + \beta_f^\top \mathbf{k}_j + \varepsilon_{ijf}$	$\mathcal{Y}_{ijf} = Y_{ijf} - \hat{\gamma}_{if}$
	ComBat	$Y_{ijf} = \alpha_f + \gamma_{if} + \beta_f^\top \mathbf{k}_j + \delta_i \varepsilon_{ijf}$	$\mathcal{Y}_{ijf} = \frac{Y_{ijf} - \hat{\alpha}_f - \hat{\beta}_f^\top \mathbf{k}_j - \hat{\gamma}_{if}}{\hat{\delta}_i} + \hat{\alpha}_f + \hat{\beta}_f^\top \mathbf{k}_j$
Site & Keep & Remove	Linear	$Y_{ijf} = \alpha_f + \gamma_{if} + \beta_f^\top \mathbf{k}_j + \zeta_f^\top \mathbf{r}_j + \varepsilon_{ijf}$	$\mathcal{Y}_{ijf} = Y_{ijf} - \hat{\gamma}_{if} - \hat{\zeta}_f^\top \mathbf{r}_j$
	ComBat	$Y_{ijf} = \alpha_f + \gamma_{if} + \beta_f^\top \mathbf{k}_j + \zeta_f^\top \mathbf{r}_j + \delta_i \varepsilon_{ijf}$	$\mathcal{Y}_{ijf} = \frac{Y_{ijf} - \hat{\alpha}_f - \hat{\beta}_f^\top \mathbf{k}_j - \hat{\zeta}_f^\top \mathbf{r}_j - \hat{\gamma}_{if}}{\hat{\delta}_i} + \hat{\alpha}_f + \hat{\beta}_f^\top \mathbf{k}_j$

Table 2: Summary of harmonization models for image features Y with the corresponding update equations for obtaining the harmonized value \mathcal{Y} . The indices are imaging site i , subject j , and feature type f . We consider linear regression and ComBat, together with different types for variables. Site is represented by γ . For biological variation that we would like to keep \mathbf{k} , we consider age and sex. For non-biological variation that we would like to remove \mathbf{r} , we consider manufacturer, field strength, and principal components. Estimated model parameters are denoted with hat.

that allows inferring multiple causal effects in the presence of unknown confounders. Our model is a non-probabilistic version of the deconfounder, which supports the theoretical connection between our approach and causal inference. We note that such an approach only helps to identify causal effects if the neurological process we are studying is believed to only affect a few brain structures. As variables for the remove vector \mathbf{r} , we are therefore considering manufacturer and magnetic field strength, together with the principal component.

6. Results

We first present results for *Name That Dataset* on harmonized features, brain age prediction, and a comparison of causal and confounded models with and without harmonization.

6.1. Name That Dataset with Harmonization

Fig. 4 shows the results of *Name That Dataset*, when the different harmonization techniques were applied. We use volume measures, healthy subjects and 70% training data for these experiments. Imaging site is used in all linear and ComBat models, further augmented with biological variables age and sex. Next, we add non-biological variables that we would like to remove, namely manufacturer (MF), magnetic field strength (MFS), and principal component (PC). As noted above, all models that include age have a linear and quadratic aging term.

While we want to lower the accuracy for predicting the dataset, we also want to keep relevant information. As a measure for the latter, we use the fact that ventricles grow with age while atrophy of the hippocampus increases. We compute the Spearman’s rank correlation of lateral ventricles and hippocampus volumes with age, to account for the non-linear relationship between the volume of brain structures and age. The objective of this experiment is

to validate whether relevant information is retained in the image-derived measurements after harmonization. If biologically relevant information is retained, we should get a positive correlation for lateral ventricles and a negative correlation for hippocampus volume. Both volumes are considered relative to the intracranial volume. A trivial solution that would impede predicting the dataset would be to set all variables to constant values, however, this would also remove all relevant information. Classification accuracies of a random forest are depicted in fig. 4A for the original brain measures, Z-Score normalized, and five different variable combinations for the linear and ComBat model.

Overall, we observe that ComBat is superior to the linear model in this experiment, which indicates the benefit of not only having an additive but also a multiplicative imaging site effect in the model. Z-Score normalization leads to a steep decrease in classification accuracy from 69% to 43%, but also the correlation with age is almost entirely removed (figs. 4B and C). Hence, not only dataset-specific information, but also relevant biological information, is removed. We observe a similar result for harmonization by residuals based only on imaging site, where ComBat yields a lower classification accuracy than the linear model (42% vs. 54%). Adding age and sex increases the classification accuracy, which is expected when adding these variables to the residuals computation. The corresponding ComBat model is closest to the model used by Fortin et al. (2018) and results in an accuracy of 50% and a change in the correlation with age of 0.04. With our extension of including additional variables to remove (MF, MFS, PC), the accuracy decreases to 45% for ComBat, while at the same time strengthening the correlation with age (0.62 and -0.38). However, whether the effect of MF and MFS is removed does not change the results for neither the linear model nor ComBat. Finally, we note that all classification accuracies are well above 37.9%, which we reported earlier for



Figure 4: A: Bars show mean dataset classification accuracy with lines indicating standard deviation for the raw measures and the different harmonization techniques. B: Spearman’s rank correlation of lateral ventricles volume with age. C: Spearman’s rank correlation of hippocampus volume with age. For the linear and ComBat models, the variables to remove r and to keep k are listed above the plots. Variables that we intend on removing are imaging site, manufacturer (MF), magnetic field strength (MFS), principal component (PC).

the dataset classification task with age and sex, indicating that the features retain dataset-specific information after harmonization.

Fig. 5 shows harmonization results for using dataset instead of imaging site as grouping variable. Since *Name the Dataset* tries to identify the dataset a scan is part of, it seems natural to harmonize with respect to dataset. The figure shows the results for Z-Score normalization, together with the linear and ComBat models that include all variables (age, sex, MF, MFS, PC). With a classification accuracy of about 50% for ComBat, this is well above the results for using imaging site. Z-Score normalization and the linear model yield accuracies of about 60%, where the former again removes almost all correlation with age. We can therefore conclude that operating on the level of imaging site is leading to better results than operating on the dataset level. This indicates a high heterogeneity within the included multi-site datasets and the necessity for modeling site-specific variations.

6.2. Results on Brain Age Prediction

In a further experiment, we evaluate brain age prediction (Franke et al., 2010; Cole and Franke, 2017; Becker et al., 2018), i.e., the prediction of a person’s age from a brain MRI scan, across datasets. The prediction of the brain age is of interest as it was demonstrated that brain age relates to cognitive aging and that the difference to the chronological age is associated to neurodegenerative diseases. A prerequisite for age prediction is sufficient data to cover the full age range of interest, which often requires combining data from multiple datasets. However, site-specific characteristics may cause unwanted variation between training and testing sets, as well as heterogeneity in the training set, which can deteriorate prediction results. Hence, we examine whether harmonization of the image features can decrease the error.

Age prediction is a multivariate regression task, for which we use random forest regression on brain volumes. We only select healthy subjects for this experiment. We

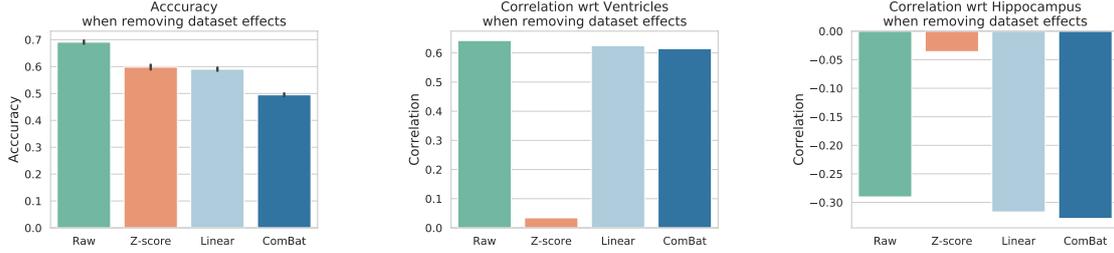


Figure 5: Dataset classification accuracy and Spearman’s rank correlation of lateral ventricles and hippocampus with age. In contrast to Fig. 4, the dataset is used as grouping variable and not the imaging site.

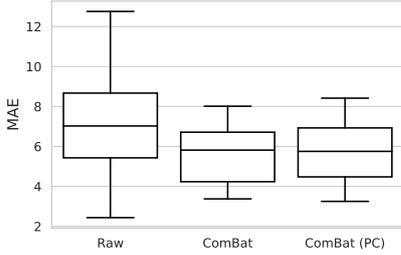


Figure 6: Boxplot of the mean absolute error (MAE) in age prediction for the raw measurements and the harmonization with ComBat and ComBat with PC. Center lines indicate the median, the boxes extend to the 25th and 75th percentiles, and the whiskers reach to the most extreme values.

use a leave-one-dataset-out scheme for the evaluation, i.e., one dataset is selected as test set and the remaining 16 datasets are used for training. Fig. 6 shows a boxplot of the mean absolute error in age prediction for the raw measurements, and the harmonized measurements with ComBat and ComBat with PC, respectively. We compute the Wilcoxon signed-rank test between the three approaches yielding $p = 0.0064$ between raw and ComBat, $p = 0.0197$ between raw and ComBatPC, and $p = 0.5171$ between ComBat and ComBatPC; the differences between the two harmonization techniques are not statistically significant, but the reduction in age prediction errors with respect to raw measures are significant.

6.3. Results on Causality with Harmonization

We estimate the causal inference model in two experiments. In the first one, we select the ADNI dataset and set Y to the Alzheimer’s Disease Assessment Scale Cognitive Subscale 13 (ADAS) score. ADAS is one of the most widely used cognitive scales and assesses the severity of cognitive symptoms of dementia (Kueper et al., 2018). Instead of directly working with diagnosis, ADAS is a continuous variable and can therefore directly be used in our causal inference framework. For the variable \mathbf{X} , we use age, age squared, sex, and the following brain volumes that have previously been associated with dementia: hippocampus, amygdala, lateral ventricles, inferior lateral ventricles, and third ventricle (Apostolova et al., 2012). The model complexities for causal and confounded models, together with the difference Δ for the raw volume measurements

is reported in Table 3. Next, we evaluate different harmonization methods for the volume measures, where we focus on ComBat. When only using site, or site in combination with age and sex, the intervals for Δ overlap, which is evidence for both working equally well. The inclusion of PC yields a higher median Δ and its interval is clearly separated from the results above, which indicates a less confounded model. Adding ADAS in the harmonization to the variables to keep (\mathbf{k}), further reduces confounding bias significantly, as is evident from non-overlapping intervals for Δ .

In the second experiment, we want to include all datasets in the study. To this end, we set Y to age and \mathbf{X} to the following brain volumes: hippocampus, lateral ventricles, amygdala, inferior lateral ventricles, third ventricle, putamen, pallidum, caudate, thalamus proper, cerebellum cortex, cerebellum white matter, and brain stem, together with sex. The results per dataset are shown in Fig. 7, where we present the normalized Δ for better comparison across datasets. Next to the median values displayed in Fig. 7, we list the range in table A.4. We report results for raw image measures and the harmonization with ComBat. As variables, we use (i) site, age, sex, (ii) site, age, sex, PC, and (iii) site, age, sex, MF, MFS, PC. For (i), which corresponds to the original version of ComBat, we observe similar values to the raw measures. After including PC, we observe a substantial increase in Δ for almost all datasets. Notably, for NKI and OASIS, the difference gets positive, which indicates a causal relationship. To investigate this further, we plot L_{ca} for NKI and OASIS in the inset in Fig. 7, which shows the large drop in complexity for the causal model after harmonization with PC. With the complexity of the confounded model changing only little, it is the change of the causal model that yields the positive Δ and therefore evidence for a causal relationship.

7. Discussion

7.1. Name That Dataset

In one of the largest studies to date, we have examined bias across 17 neuroimaging datasets. Our experiments for *Name That Dataset* demonstrated that the dataset can be predicted with more than 70% accuracy. In comparison, a classifier that operates on age and sex stays below

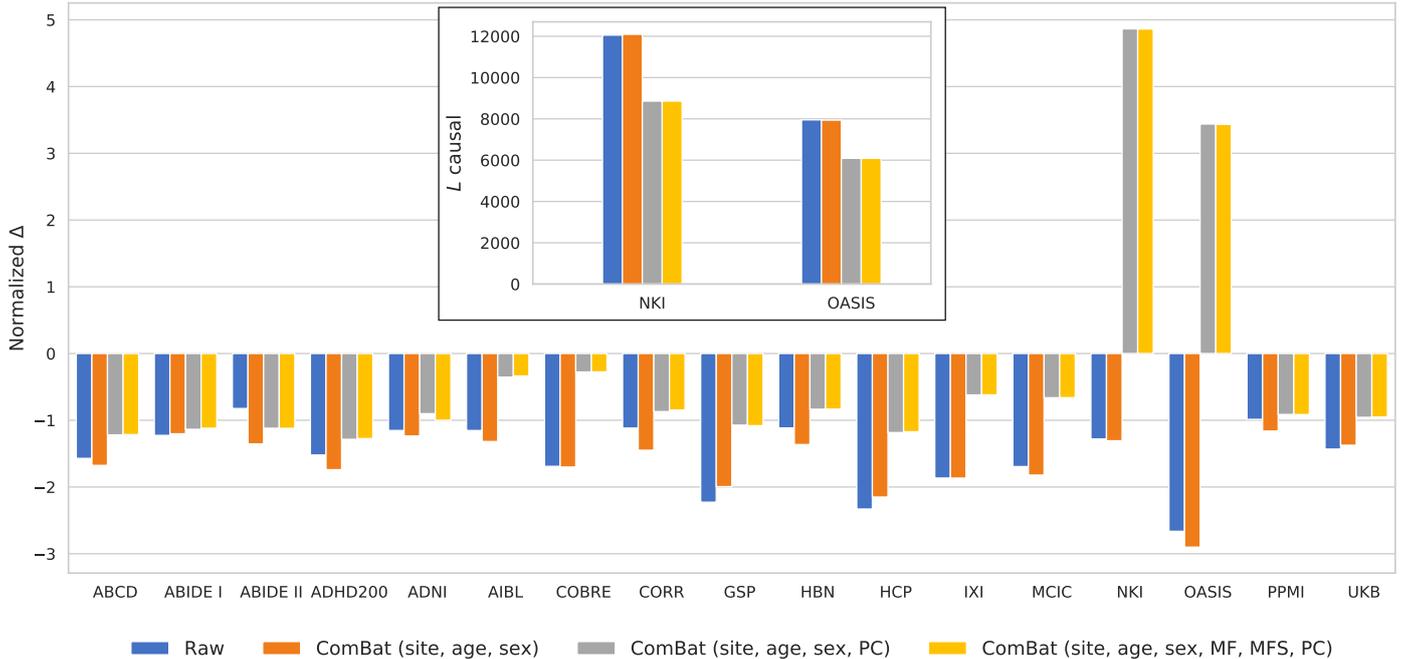


Figure 7: Results from the causal inference model across all datasets, where the median normalized Δ is plotted. Results are shown for the raw image measurements and the ComBat harmonization with different variable combinations. The inset plots L_{ca} for NKI and OASIS. Minimum and maximum normalized Δ values across 10 repetitions are available in table A.4.

Table 3: Results from the causal inference model with the variable Y set to ADAS for the raw image measurements and the application of harmonization with different variables. Median values are reported for L_{ca} , L_{co} , and Δ , and the minimum and maximum for Δ .

	L_{ca} (Median)	L_{co} (Median)	Δ (Median)	Δ (Range)
Raw	20,238	18,782	-1,456	[-1,451; -1,461]
Site	20,221	18,789	-1,432	[-1,428; -1,435]
Site, Age, Sex	20,203	18,772	-1,430	[-1,429; -1,437]
Site, Age, Sex, PC	20,283	19,097	-1,186	[-1,184; -1,187]
Site, Age, Sex, PC, ADAS	20,185	19,362	-823	[-821; -827]

38%. This is evidence that image features contain, next to subject-specific information, also rich dataset-specific information. While we are not aware of previous work on predicting the dataset from neuroimaging data, prior studies have reported the high intra-subject variability across scanners, no matter whether FreeSurfer (Jovicich et al., 2009), FSL (Nugent et al., 2013; Suckling et al., 2012), or BrainVisa (Shokouhi et al., 2011) were used for segmentation, which is also pointing to scanner-specific information in the data.

Instead of predicting the dataset, a similar experiment would be to predict manufacturer or magnetic field strength of a scan. We believe that dataset prediction offers two advantages. First, there is a severe class imbalance with almost 80% of scans acquired on a Siemens scanner and 84% on 3.0T, yielding already high classification accuracies with naïve approaches. Second, there are more non-

biological effects than just manufacturer and field strength that impact a scan; by using the dataset as label, we can lump them all together, without explicitly knowing them. However, we also need to account for biological variation between datasets, which could be considered as a lower bound for the classification accuracy, as we would like to retain this information. In our experiments, we used a classification model based on the variables age and sex for this purpose. Another alternative to *Name That Dataset* would be to *Name That Imaging Site*. With 210 sites in our data, this would substantially increase the complexity of the classification task and complicate the interpretation of results with the confusion matrix. In addition, there is a wide variation in the size of sites, ranging from a few subjects per site to more than 1,000 for HCP. Hence, *Name That Dataset* is the most reasonable choice in our experiments, but operating on sites could be an attractive

alternative in other situations.

7.2. Dataset Harmonization

We have investigated several techniques for dataset harmonization and evaluated whether they can reduce the accuracy of predicting the dataset or whether they can yield a causal relationship. Harmonization is a trade-off between removing unwanted variation, likely associated to scanner, and keeping relevant information about the research question. Hence, we need to be cautious when harmonizing image features. As shown in our results for *Name That Dataset*, Z-score normalization and the harmonization with imaging site can only drastically decrease the classification accuracy. However, the low correlation of age with hippocampus and ventricles volumes after harmonization illustrates that all the biologically relevant information has been removed too. Adding more supervision to the process is therefore necessary, which we achieved by adding variables whose association we want to keep in the data (e.g. age and sex). Results from our correlation analyses with age show the effectiveness of such an approach, but also that the dataset prediction accuracy increases compared to only using site.

As a novel contribution to harmonization, we proposed to add variables to the model, whose influence we would like to remove. In particular, we proposed to compute principal components across all image features, and add them as variables to remove. This approach is inspired from GWAS, where principal components across genetic markers are added to a model of association to account for stratification of the population. In contrast, we use principal components in a regression model to harmonize imaging data by computing the residuals. Next to PC, our technique can also be used to explicitly remove the effect of measured variables; we studied scanner manufacturer and magnetic field strength. In combination with imaging site as grouping variable, their removal had only small impact. This is most likely due to manufacturer and field strength being constant within an imaging site, so that the information provided by these variables is already accounted for by imaging site. Manufacturer and field strength have a stronger influence on the results when dataset is used as grouping variable.

The prediction of the brain age is a typical example, where the combination of multiple dataset is required to cover a wide age range in the training dataset. The results of our leave-one-dataset-out experiment demonstrated that harmonization yields a significant reduction in the age prediction error. As age prediction has been proposed as a diagnostic biomarker for many diseases, like AD, HIV, and schizophrenia (Cole and Franke, 2017), accounting for dataset bias may further increase its sensitivity.

For the computation of residuals in the harmonization, we worked with linear regression models and ComBat, where ComBat has an additional site-specific scaling factor and uses empirical Bayes for inference. Throughout all our

experiments, we noticed advantages of ComBat in comparison to the linear model. This confirms previous results from Fortin et al. (2018) for harmonizing cortical thickness measures with ComBat also for the harmonization of volume measurements. We have focused on harmonization approaches that operate on image-derived features, while previous approaches harmonized image intensity values to make them comparable across studies, including histogram matching (Nyúl et al., 2000), WhiteStripe (Shinohara et al., 2014), and RAVEL (Fortin et al., 2016). As an alternative to accounting for population stratification with PCs added to the model, mixed linear models have been proposed for genetic association studies (Yang et al., 2014), where the random effect accounts for stratification. We could adopt a similar approach to model unknown effects in image features. However, since random effects are not explicitly estimated, correcting the image features in a harmonization step is not directly possible, but rather an integrated model would be required.

7.3. Causal Inference for Identifying Confounding

Confounding bias can lead to spurious correlations and therefore wrong conclusions about cause and effect. The assumption of causal sufficiency, i.e., knowing all confounding variables, is often violated in practice. Hence, we presented a causal inference framework that considers *unknown confounders*. In contrast, prior work in neuroimaging (Dukart et al., 2011; Linn et al., 2016; Rao et al., 2017) mainly considered the known confounders age and sex. We believe that our approach can present an interesting new analysis for neuroimaging data, e.g., for evaluating methods for harmonization or bias reduction.

Our results indicate that most relationships that we investigated were confounded rather than causal. This is unsurprising, given the inherent complexity of neuroimaging studies related to scanning, image analysis, and study design. On the positive side, techniques for harmonization can reduce confounding and, in some cases, even yield causal relationships. For both causal experiments, we observed that accounting for unknown confounders in ComBat via principal components is preferred. In addition, the experiment on the ADNI dataset revealed that including the target variable (ADAS) as variable to keep, results in an improved harmonization that reduces confounding significantly. However, the relationship remains non-causal ($\Delta < 0$). We believe that morphology alone may be insufficient, and that additional, multi-modal measures would be necessary to establish a causal model for ADAS.

As mentioned earlier, inferring causality from observational data is challenging and only feasible under certain assumptions as stated in section 4. These assumptions heavily rely on domain knowledge and may not hold for other applications. For instance, we assumed a linear relationship between continuous-valued variables, which suggested the use of linear regression as causal model and probabilistic PCA as confounded model. If non-linearity

is required, more advanced models such as Gaussian Process regression and probabilistic non-linear PCA could be used (Rasmussen and Williams, 2006; Lawrence, 2005). Similarly, if data comprises discrete variables, appropriate models need to be selected, such as Poisson regression as casual model and Gamma-Poisson factorization as confounded model (Canny, 2004). All of these models would fit the setting depicted in Fig. 3 and allow estimation of Δ as described in section 4. In addition to relying on untestable assumptions, causal inference from observational data is challenging, because we cannot know what the true causal effects are nor their effect size, which makes a quantitative evaluation impossible. As such, results need to be interpreted with caution and in light of the incorporated prior knowledge. Yet, the ability to assess confounding in the model can be very helpful, not just for the development of harmonization tools, but more generally for inferring knowledge from neuroimaging studies.

8. Conclusion

Bias is a complex and challenging topic in neuroimaging that will become more prevalent in the future with the translation of results in the clinic and the surge of mega-analyses. Based on data with more than 35,000 individuals, we have demonstrated that simply pooling scans from distinct studies can introduce substantial bias that would be passed on to a machine learning model trained on the pooled data. First, we showed that it is possible to correctly identify the dataset that a scan is part of with more than 70% accuracy. Second, we introduced a novel approach for differentiating causal from confounded relationships based on causal inference. Importantly, the confounder was modeled as unknown variable, which is helpful in complex neuroimaging studies, where the assumption of causal sufficiency is challenging to fulfill. Third, we studied multi-site harmonization techniques and evaluated their effectiveness of reducing bias. We have proposed a harmonization method that extends ComBat for the inclusion of additional variables to remove, where the integration of principal components, to capture generic variation, led to the best results.

Overall, we believe that the growing amount of medical images necessitates novel methods for handling bias in datasets and image-derived features. As bias is in its core a causal concept, methods from the growing field of causal inference may be particularly promising to yield new insights.

Acknowledgements

This research was supported by the Bavarian State Ministry of Science and the Arts and coordinated by the Bavarian Research Institute for Digital Transformation (bidt). The authors gratefully acknowledge the Leibniz Supercomputing Centre for providing computing time on its Linux-Cluster.

We acknowledge the following initiatives for providing data.

ABCD: Data used in the preparation of this article were obtained from the Adolescent Brain Cognitive Development (ABCD) Study (<https://abcdstudy.org>), held in the NIMH Data Archive (NDA). This is a multisite, longitudinal study designed to recruit more than 10,000 children aged 9-10 years and follow them over 10 years into early adulthood. The ABCD Study is supported by the National Institutes of Health and additional federal partners under award numbers U01DA041022, U01DA041028, U01DA041048, U01DA041089, U01DA041106, U01DA041117, U01DA041120, U01DA041134, U01DA041148, U01DA041156, U01DA041174, U24DA041025, U01DA041093, U24DA041123, and U24DA041147. A full list of supporters is available at <https://abcdstudy.org/nih-collaborators>. A listing of participating sites and a complete listing of the study investigators can be found at <https://abcdstudy.org/principal-investigators.html>. ABCD consortium investigators designed and implemented the study and/or provided data but did not necessarily participate in analysis or writing of this report. This manuscript reflects the views of the authors and may not reflect the opinions or views of the NIH or ABCD consortium investigators.

ADNI: Data collection and sharing for this project was funded by the Alzheimer’s Disease Neuroimaging Initiative (ADNI) (National Institutes of Health Grant U01 AG024904) and DOD ADNI (Department of Defense award number W81XWH-12-2-0012). ADNI is funded by the National Institute on Aging, the National Institute of Biomedical Imaging and Bioengineering, and through generous contributions from the following: Alzheimer’s Association; Alzheimer’s Drug Discovery Foundation; Araclon Biotech; BioClinica, Inc.; Biogen Idec Inc.; Bristol-Myers Squibb Company; Eisai Inc.; Elan Pharmaceuticals, Inc.; Eli Lilly and Company; EuroImmun; F. Hoffmann-La Roche Ltd and its affiliated company Genentech, Inc.; Fujirebio; GE Healthcare; ; IXICO Ltd.; Janssen Alzheimer Immunotherapy Research & Development, LLC.; Johnson & Johnson Pharmaceutical Research & Development LLC.; Medpace, Inc.; Merck & Co., Inc.; Meso Scale Diagnostics, LLC.; NeuroRx Research; Neurotrack Technologies; Novartis Pharmaceuticals Corporation; Pfizer Inc.; Piramal Imaging; Servier; Synarc Inc.; and Takeda Pharmaceutical Company. The Canadian Institutes of Health Research is providing funds to support ADNI clinical sites in Canada. Private sector contributions are facilitated by the Foundation for the National Institutes of Health (www.fnih.org). The grantee organization is the Northern California Institute for Research and Education, and the study is coordinated by the Alzheimer’s Disease Cooperative Study at the University of California, San Diego. ADNI data are disseminated by the Laboratory for Neuro Imaging at the University of Southern California.

COBRE: Data was downloaded from the COllaborative Informatics and Neuroimaging Suite Data Exchange

tool (COINS; <http://coins.mrn.org/dx>) and data collection was performed at the Mind Research Network, and funded by a Center of Biomedical Research Excellence (COBRE) grant 5P20RR021938/P20GM103472 from the NIH to Dr. Vince Calhoun.

MIND: Data used in the preparation of this work were obtained from the Mind Clinical Imaging Consortium database through the Mind Research Network (www.mrn.org). The MCIC project was supported by the Department of Energy under Award Number DE-FG02-08ER64581. MCIC is the result of efforts of co-investigators from University of Iowa, University of Minnesota, University of New Mexico, Massachusetts General Hospital.

GSP: Data were provided in part by the Brain Genomics Superstruct Project of Harvard University and the Massachusetts General Hospital, (Principal Investigators: Randy Buckner, Joshua Roffman, and Jordan Smoller), with support from the Center for Brain Science Neuroinformatics Research Group, the Athinoula A. Martinos Center for Biomedical Imaging, and the Center for Human Genetic Research. 20 individual investigators at Harvard and MGH generously contributed data to the overall project.

HCP: Data were provided in part by the Human Connectome Project, WU-Minn Consortium (Principal Investigators: David Van Essen and Kamil Ugurbil; 1U54MH091657) funded by the 16 NIH Institutes and Centers that support the NIH Blueprint for Neuroscience Research; and by the McDonnell Center for Systems Neuroscience at Washington University

OASIS: Data were provided in part by OASIS Principal Investigators: D. Marcus, R. Buckner, J. Csernansky J. Morris; P50 AG05681, P01 AG03991, P01 AG026276, R01 AG021910, P20 MH071616, U24 RR021382.

PPMI: Data used in the preparation of this article were obtained from the Parkinson’s Progression Markers Initiative (PPMI) database (www.ppmi-info.org/data). For up-to-date information on the study, visit www.ppmi-info.org. PPMI a public-private partnership is funded by the Michael J. Fox Foundation for Parkinson’s Research and funding partners, including [list the full names of all of the PPMI funding partners found at www.ppmi-info.org/fundingpartners].

UKB: This research has been conducted using the UK Biobank Resource under the application number 34479.

9. References

References

Alexander, L. M., Escalera, J., et al., 2017. An open resource for transdiagnostic research in pediatric mental health and learning disorders. *bioRxiv*, 149369.

Apostolova, L. G., Green, A. E., Babakchianian, S., Hwang, K. S., Chou, Y.-Y., Toga, A. W., Thompson, P. M., 2012. Hippocampal atrophy and ventricular enlargement in normal aging, mild cognitive impairment and alzheimer’s disease. *Alzheimer disease and associated disorders* 26 (1), 17.

Becker, B. G., Klein, T., Wachinger, C., 2018. Gaussian process uncertainty in age estimation as a measure of brain abnormality. *NeuroImage*.

Buckner, R., Hollinshead, M., Holmes, A., Brohawn, D., Fagerness, J., O’Keefe, T., Roffman, J., 2012. The brain genomics superstruct project. HDN.

Canny, J., 2004. Gap: A factor model for discrete data. In: *Proceedings of the 27th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*. pp. 122–129.

Carpenter, B., Gelman, A., Hoffman, M. D., Lee, D., Goodrich, B., Betancourt, M., Brubaker, M., Guo, J., Li, P., Riddell, A., 2017. Stan: A probabilistic programming language. *Journal of Statistical Software* 76 (1).

Casey, B., Cannonier, T., Conley, M. I., Cohen, A. O., Barch, D. M., Heitzeg, M. M., Soules, M. E., Teslovich, T., Dellarco, D. V., Garavan, H., et al., 2018. The adolescent brain cognitive development (abcd) study: imaging acquisition across 21 sites. *Developmental cognitive neuroscience* 32, 43–54.

Cole, J. H., Franke, K., 2017. Predicting age using neuroimaging: innovative brain ageing biomarkers. *Trends in neurosciences* 40 (12), 681–690.

Di Martino, A., O’connor, D., Chen, B., Alaerts, K., Anderson, J. S., Assaf, M., Balsters, J. H., Baxter, L., Beggiano, A., Bernaerts, S., et al., 2017. Enhancing studies of the connectome in autism using the autism brain imaging data exchange ii. *Scientific data* 4 (1), 1–15.

Di Martino, A., Yan, C., et al., 2014. The autism brain imaging data exchange: towards a large-scale evaluation of the intrinsic brain architecture in autism. *Molecular psychiatry* 19 (6), 659–667.

Dukart, J., Schroeter, M. L., Mueller, K., 2011. Age correction in dementia—matching to a healthy brain. *PloS one* 6 (7), e22193.

Ellis, K., Bush, A., Darby, D., et al., 2009. The australian imaging, biomarkers and lifestyle (aibl) study of aging. *International Psychogeriatrics* 21 (04), 672–687.

Fischl, B., Salat, D. H., et al., 2002. Whole brain segmentation: automated labeling of neuroanatomical structures in the human brain. *Neuron* 33 (3), 341–355.

Fortin, J.-P., Cullen, N., et al., 2018. Harmonization of cortical thickness measurements across scanners and sites. *Neuroimage* 167, 104–120.

Fortin, J.-P., Sweeney, E. M., Muschelli, J., Crainiceanu, C. M., Shinohara, R. T., Alzheimer’s Disease Neuroimaging Initiative, et al., 2016. Removing inter-subject technical variability in magnetic resonance imaging studies. *Neuroimage* 132, 198–212.

Franke, K., Ziegler, G., Klöppel, S., Gaser, C., Initiative, A. D. N., et al., 2010. Estimating the age of healthy subjects from t1-weighted mri scans using kernel methods: exploring the influence of various parameters. *Neuroimage* 50 (3), 883–892.

Frühwirth-Schnatter, S., 2004. Estimating marginal likelihoods for mixture and Markov switching models using bridge sampling techniques. *The Econometrics Journal* 7 (1), 143–167.

Gollub, R. L., Shoemaker, J., King, M., White, T., Ehrlich, S., Sponheim, S., Clark, V., Turner, J., Mueller, B., Magnotta, V., et al., 2013. The mcic collection: a shared repository of multi-modal, multi-site brain image data from a clinical investigation of schizophrenia. *Neuroinformatics* 11 (3), 367–388.

Gronau, Q., Singmann, H., Wagenmakers, E.-J., 2020. bridgesampling: An r package for estimating normalizing constants. *Journal of Statistical Software* 92 (10), 1–29.

Grünwald, P. D., 2007. *The Minimum Description Length Principle*. MIT press.

Guadalupe, T., Mathias, S. R., Theo, G., et al., 2017. Human subcortical brain asymmetries in 15,847 people worldwide reveal effects of age and sex. *Brain imaging and behavior* 11 (5), 1497–1514.

Han, X., Fischl, B., 2007. Atlas renormalization for improved brain mr image segmentation across scanner platforms. *IEEE TMI* 26 (4), 479–486.

Hernan, M., Robins, J., 2020. *Causal Inference: What If*. Boca Raton.

Hoffman, M. D., Gelman, A., 2014. The no-u-turn sampler: Adaptively setting path lengths in hamiltonian monte carlo. *Journal of*

- Machine Learning Research 15, 1593–1623.
- Jack, C. R., Bernstein, M. A., Fox, N. C., Thompson, P., et al., 2008. The alzheimer’s disease neuroimaging initiative (adni): Mri methods. *Journal of magnetic resonance imaging* 27 (4), 685–691.
- Janzing, D., Scholkopf, B., 2010. Causal inference using the algorithmic markov condition. *IEEE Transactions on Information Theory* 56 (10), 5168–5194.
- Johnson, W. E., Li, C., Rabinovic, A., 2007. Adjusting batch effects in microarray expression data using empirical bayes methods. *Biostatistics* 8 (1), 118–127.
- Jovicich, J., Czanner, S., Han, X., Salat, D., van der Kouwe, A., Quinn, B., Pacheco, J., Albert, M., Killiany, R., Blacker, D., et al., 2009. Mri-derived measurements of human subcortical, ventricular and intracranial brain volumes: reliability effects of scan sessions, acquisition sequences, data analyses, scanner upgrade, scanner vendors and field strengths. *Neuroimage* 46 (1), 177–192.
- Kaltenpoth, D., Vreeken, J., 2019. We are not your real parents: Telling causal from confounded by mdl. In: *SIAM International Conference on Data Mining*.
- Krugel, F., Turner, J., Muftuler, L. T., 2010. Impact of scanner hardware and imaging protocol on image quality and compartment volume precision in the adni cohort. *Neuroimage* 49 (3), 2123–2133.
- Kueper, J. K., Speechley, M., Montero-Odasso, M., 2018. The alzheimer’s disease assessment scale–cognitive subscale (adascog): modifications and responsiveness in pre-dementia populations. a narrative review. *Journal of Alzheimer’s Disease* 63 (2), 423–444.
- Lawrence, N., 2005. Probabilistic Non-linear Principal Component Analysis with Gaussian Process Latent Variable Models. *Journal of Machine Learning Research* 6, 1783–1816.
- Liaw, A., Wiener, M., et al., 2002. Classification and regression by randomforest. *R news* 2 (3), 18–22.
- Linn, K. A., Gaonkar, B., Doshi, J., Davatzikos, C., Shinohara, R. T., 2016. Addressing confounding in predictive models with an application to neuroimaging. *The international journal of biostatistics* 12 (1), 31–44.
- Marcus, D. S., Wang, T. H., Parker, J., Csernansky, J. G., Morris, J. C., Buckner, R. L., 2007. Open access series of imaging studies (oasis): cross-sectional mri data in young, middle aged, nondemented, and demented older adults. *J. Cognitive Neurosci.* 19 (9), 1498–1507.
- Marek, K., Jennings, D., Lasch, S., Siderowf, A., Tanner, C., Simuni, T., Coffey, C., Kieburz, K., Flagg, E., Chowdhury, S., et al., 2011. The parkinson progression marker initiative (ppmi). *Progress in neurobiology* 95 (4), 629–635.
- Mayer, A., Ruhl, D., Merideth, F., Ling, J., Hanlon, F., Bustillo, J., Cañive, J., 2013. Functional imaging of the hemodynamic sensory gating response in schizophrenia. *Human brain mapping* 34 (9), 2302–2312.
- Milham, M. P., Fair, D., Mennes, M., Mostofsky, S. H., et al., 2012. The adhd-200 consortium: a model to advance the translational potential of neuroimaging in clinical neuroscience. *Frontiers in systems neuroscience* 6, 62.
- Miller, K. L., Alfaro-Almagro, F., Bangarter, N. K., Thomas, D. L., Yacoub, E., Xu, J., Bartsch, A. J., Jbabdi, S., Sotiropoulos, S. N., Andersson, J. L., et al., 2016. Multimodal population brain imaging in the uk biobank prospective epidemiological study. *Nature neuroscience* 19 (11), 1523.
- Nooner, K. B., et al., 2012. The nki-rockland sample: a model for accelerating the pace of discovery science in psychiatry. *Frontiers in neuroscience* 6.
- Nugent, A. C., Luckenbaugh, D. A., Wood, S. E., Bogers, W., Zarate, C. A., Drevets, W. C., 2013. Automated subcortical segmentation using first: test–retest reliability, interscanner reliability, and comparison to manual segmentation. *Human brain mapping* 34 (9), 2313–2329.
- Nyúl, L. G., Udupa, J. K., Zhang, X., 2000. New variants of a method of mri scale standardization. *IEEE transactions on medical imaging* 19 (2), 143–150.
- Pearl, J., 2000. *Causality: Models, reasoning, and inference*. Cambridge University Press.
- Price, A. L., Patterson, N. J., Plenge, R. M., Weinblatt, M. E., Shadick, N. A., Reich, D., 2006. Principal components analysis corrects for stratification in genome-wide association studies. *Nature genetics* 38 (8), 904.
- Rao, A., Monteiro, J. M., Mourao-Miranda, J., 2017. Predictive modelling using neuroimaging data in the presence of confounds. *NeuroImage* 150, 23–49.
- Rasmussen, C., Williams, C., 2006. *Gaussian processes for machine learning*. MIT Press.
- Scahill, R. I., Frost, C., Jenkins, R., Whitwell, J. L., Rossor, M. N., Fox, N. C., 2003. A longitudinal study of brain volume changes in normal aging using serial registered magnetic resonance imaging. *Archives of neurology* 60 (7), 989–994.
- Shinohara, R. T., Sweeney, E. M., Goldsmith, J., Shiee, N., Maateen, F. J., Calabresi, P. A., Jarso, S., Pham, D. L., Reich, D. S., Crainiceanu, C. M., et al., 2014. Statistical normalization techniques for magnetic resonance imaging. *NeuroImage: Clinical* 6, 9–19.
- Shokouhi, M., Barnes, A., Suckling, J., Moorhead, T. W., Brennan, D., Job, D., Lymer, K., Dazzan, P., Marques, T. R., MacKay, C., et al., 2011. Assessment of the impact of the scanner-related factors on brain morphometry analysis with brainvisa. *BMC Medical Imaging* 11 (1), 23.
- Smith, S. M., Nichols, T. E., 2018. Statistical challenges in “big data” human neuroimaging. *Neuron* 97 (2), 263–268.
- Stan Development Team, 2020. *RStan: the R interface to Stan*. R package version 2.19.3. URL <http://mc-stan.org/>
- Suckling, J., Barnes, A., Job, D., Brennan, D., Lymer, K., Dazzan, P., Marques, T. R., MacKay, C., McKie, S., Williams, S. R., et al., 2012. The neuro/psygrid calibration experiment: identifying sources of variance and bias in multicenter mri studies. *Human brain mapping* 33 (2), 373–386.
- Tipping, M. E., Bishop, C. M., Aug. 1999. Probabilistic Principal Component Analysis. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 61 (3), 611–622.
- Torralba, A., Efros, A. A., 2011. Unbiased look at dataset bias. In: *Computer Vision and Pattern Recognition (CVPR)*. pp. 1521–1528.
- Van Essen, D. C., Smith, S. M., Barch, D. M., Behrens, T., Yacoub, E., Ugurbil, K., Consortium, W.-M. H., et al., 2013. The wu-minn human connectome project: an overview. *Neuroimage* 80, 62–79.
- Wachinger, C., Becker, B. G., Rieckmann, A., Pölsterl, S., 2019. Quantifying confounding bias in neuroimaging datasets with causal inference. In: *Medical Image Computing and Computer-Assisted Intervention–MICCAI*. Springer.
- Wachinger, C., Reuter, M., 2016. Domain adaptation for alzheimer’s disease diagnostics. *Neuroimage* 139, 470–479.
- Walhovd, K. B., Westlye, L. T., Amlien, I., Espeseth, T., Reinvang, I., Raz, N., Agartz, I., Salat, D. H., Greve, D. N., Fischl, B., et al., 2011. Consistent neuroanatomical age-related volume differences across multiple samples. *Neurobiology of aging* 32 (5), 916–932.
- Wang, Y., Blei, D. M., 2019. The blessings of multiple causes. *Journal of the American Statistical Association* 114 (528), 1574–1596.
- Yang, J., Zaitlen, N., Goddard, M., Visscher, P., Price, A., 2014. Mixed model association methods: advantages and pitfalls. *Nat. Genet* 46 (2), 100–106.
- Zuo, X.-N., Anderson, J. S., Bellec, P., et al., 2014. An open science resource for establishing reliability and reproducibility in functional connectomics. *Scientific data* 1, 140049.

Appendix A. Additional Results

Table A.4: Median, minimum, and maximum normalized Δ across 10 repetitions from the causal inference model across all datasets (cf. Fig. 7).

Dataset	Raw	ComBat (site, age, sex)	ComBat (site, age, sex, PC)	ComBat (site, age, sex, MF, MFS, PC)
ABCD	-1.571 [-1.572; -1.570]	-1.676 [-1.677; -1.675]	-1.219 [-1.220; -1.218]	-1.214 [-1.216; -1.213]
ABIDE I	-1.226 [-1.229; -1.224]	-1.202 [-1.205; -1.201]	-1.135 [-1.139; -1.133]	-1.119 [-1.121; -1.117]
ABIDE II	-0.824 [-0.825; -0.823]	-1.354 [-1.356; -1.352]	-1.119 [-1.125; -1.118]	-1.120 [-1.121; -1.117]
ADHD200	-1.520 [-1.527; -1.519]	-1.740 [-1.742; -1.738]	-1.285 [-1.293; -1.282]	-1.274 [-1.276; -1.273]
ADNI	-1.153 [-1.155; -1.152]	-1.235 [-1.238; -1.233]	-0.901 [-0.902; -0.898]	-1.000 [-1.003; -0.998]
AIBL	-1.154 [-1.154; -1.154]	-1.318 [-1.318; -1.318]	-0.354 [-0.354; -0.353]	-0.336 [-0.337; -0.336]
COBRE	-1.691 [-1.691; -1.690]	-1.701 [-1.701; -1.700]	-0.276 [-0.277; -0.276]	-0.275 [-0.275; -0.274]
CORR	-1.117 [-1.121; -1.116]	-1.449 [-1.451; -1.446]	-0.870 [-0.873; -0.868]	-0.847 [-0.851; -0.845]
GSP	-2.228 [-2.232; -2.225]	-1.995 [-1.998; -1.993]	-1.073 [-1.079; -1.071]	-1.081 [-1.085; -1.079]
HBN	-1.115 [-1.115; -1.114]	-1.363 [-1.365; -1.362]	-0.834 [-0.836; -0.832]	-0.833 [-0.838; -0.832]
HCP	-2.330 [-2.337; -2.327]	-2.150 [-2.156; -2.148]	-1.184 [-1.186; -1.182]	-1.172 [-1.174; -1.170]
IXI	-1.866 [-1.867; -1.865]	-1.866 [-1.868; -1.863]	-0.620 [-0.621; -0.619]	-0.619 [-0.621; -0.618]
MCIC	-1.694 [-1.694; -1.694]	-1.821 [-1.821; -1.821]	-0.662 [-0.662; -0.662]	-0.661 [-0.661; -0.660]
NKI	-1.280 [-1.281; -1.279]	-1.307 [-1.308; -1.307]	4.864 [4.862; 4.866]	4.863 [4.861; 4.864]
OASIS	-2.664 [-2.664; -2.664]	-2.901 [-2.901; -2.901]	3.439 [3.439; 3.439]	3.433 [3.433; 3.433]
PPMI	-0.986 [-0.987; -0.986]	-1.161 [-1.162; -1.161]	-0.913 [-0.913; -0.912]	-0.915 [-0.915; -0.914]
UKB	-1.432 [-1.433; -1.432]	-1.374 [-1.374; -1.372]	-0.955 [-0.956; -0.955]	-0.952 [-0.953; -0.951]