

Estimation of Causal Effects in the Presence of Unobserved Confounding in the Alzheimer’s Continuum

Sebastian Pölsterl and Christian Wachinger

Artificial Intelligence in Medical Imaging (AI-Med),
Department of Child and Adolescent Psychiatry,
Ludwig-Maximilians-Universität, Munich, Germany
{sebastian.poelsterl,christian.wachinger}@med.uni-muenchen.de

Abstract. Studying the relationship between neuroanatomy and cognitive decline due to Alzheimer’s has been a major research focus in the last decade. However, to infer cause-effect relationships rather than simple associations from observational data, we need to (i) express the causal relationships leading to cognitive decline in a graphical model, and (ii) ensure the causal effect of interest is identifiable from the collected data. We derive a causal graph from the current clinical knowledge on cause and effect in the Alzheimer’s disease continuum, and show that identifiability of the causal effect requires all confounders to be known and measured. However, in complex neuroimaging studies, we neither know all potential confounders nor do we have data on them. To alleviate this requirement, we leverage the dependencies among multiple causes by deriving a substitute confounder via a probabilistic latent factor model. In our theoretical analysis, we prove that using the substitute confounder enables identifiability of the causal effect of neuroanatomy on cognition. We quantitatively evaluate the effectiveness of our approach on semi-synthetic data, where we know the true causal effects, and illustrate its use on real data on the Alzheimer’s disease continuum, where it reveals important causes that otherwise would have been missed.

1 Introduction

The last decade saw an unprecedented increase in large multi-site neuroimaging studies, which opens the possibility of identifying disease predictors with low-effect sizes. However, one major obstacle to fully utilize this data is confounding. The relationship between a measurement and an outcome is confounded if the observed association is only due to a third latent random variable, but there is no direct causal link between the measurement and outcome. If confounding is ignored, investigators will likely make erroneous conclusions, because the observed data distribution is compatible with many – potentially contradictory – causal explanations, leaving us with no way to differentiate between the true and false effect on the basis of data. In this case, the causal effect is unidentifiable [24].

It is important to remember that what is, or is not, regarded as a confounding variable is relative and depends on the goal of the study. For instance, age is often considered a confounder when studying Alzheimer’s disease (AD), but if the focus is age-related cognitive decline in a healthy population, age is not considered a confounder [20]. Therefore, it is vital to state the causal question being studied.

Causal inference addresses confounding in a principal manner and allows us to determine which cause-effect relationships can be identified from a given dataset. To infer causal effects from observational data, we need to rely on expert knowledge and untestable assumptions about the data-generating process to build the causal graph linking causes, outcome, and other variables [24]. One essential assumption to estimate causal effects from observational data is that of *no unmeasured confounder* [24]. Usually, we can only identify causal effects if we know and recorded all confounders. However, analyses across 17 neuroimaging studies revealed that considerable bias remains in volume and thickness measurement after adjusting for age, gender, and the type of MRI scanner [33]. Another study on confounders in UK Biobank brain imaging data identified hundreds of potential confounders just related to the acquisition process researchers would need to account for [1]. These results suggest that all factors contributing to confounding in neuroimaging are not yet fully understood, and hence the premise of *no unmeasured confounder* is most likely going to be violated.

In this paper, we focus on the problem of estimating causal effects of neuroanatomical measures on cognitive decline due to Alzheimer’s in the presence of *unobserved confounders*. To make this feasible, we derive a causal graph from domain knowledge on the Alzheimer’s disease continuum to capture known disease-specific relationships. While causal affects are generally unidentifiable in the presence of unobserved confounding, we will illustrate that we can leverage the dependencies among multiple causes to estimate a latent substitute confounder via a Bayesian probabilistic factor model. In our experiments, we quantitatively demonstrate the effectiveness of our approach on semi-synthetic data, where we know the true causal effects, and illustrate its use on real data on the Alzheimer’s disease continuum, where our analyses reveal important causes of cognitive function that would otherwise have been missed.

Related Work. Despite the importance of this topic, there has been little prior work on causal inference for estimating causal effects in neuroimaging. In contrast to our approach, most of the previous works assume that all confounding variables are known and have been measured. The most common approach for confounding adjustment is regress-out. In regress-out, the original measurement (e.g. volume or thickness) is replaced by the residual of a regression model fitted to estimate the original value from the confounding variables. In [6], the authors use linear regression to account for age, which has been extended to additionally account for gender in [17]. A Gaussian Process (GP) model has been proposed in [18] to adjust for age, total intracranial volume, sex, and MRI scanner. Fortin et al. [10] proposed a linear mixed effects model to account for systematic differences across imaging sites. This model has been extended in [33] to account for both observed and unobserved confounders. In [29], linear regression is used to regress-out the

effect of total brain volume. In addition to the regress-out approach, analyses can be adjusted for confounders by computing instance weights, that are used in a downstream classification or regression model to obtain a pseudo-population that is approximately balanced with respect to the confounders [19,25]. A weighted support vector machine to adjust for age is proposed in [19]. In [25], weighted GP regression is proposed to account for gender and imaging site effects. We note that none of the work above studied whether causal effects can actually be identified from observed data using the theory of causal inference.

2 Methods

Causal inference from observational data comprises multiple steps, (i) defining the causal question and its associated causal graph, (ii) determining under which conditions the question can be answered from real world data, and (iii) estimating the causal effects via modelling. We denote random variables with uppercase letters and specific values taken by the corresponding variables with lowercase letters. We distinguish between real-valued subcortical volume ($X_1^v, \dots, X_{D_1}^v$) and cortical thickness ($X_1^t, \dots, X_{D_2}^t$) measurements. We denote by X_1, \dots, X_D all measurements, irrespective of their type ($D = D_1 + D_2$). Next, we will specify our causal question and determine when the causal effect of a subset of measurements on an outcome is identifiable using Pearl’s do-calculus [24].

2.1 The Causal Question and Its Associated Graph

Causal Question. What is the average causal effect of increasing/decreasing the volume or thickness of a subset of neuroanatomical structures on the Alzheimer’s Disease Assessment Scale Cognitive Subscale 13 score (ADAS; [23]) in patients with an Alzheimer’s pathologic change [13]?

The gold standard to answer this question would be a randomized experiment, where subjects’ volumes and thicknesses are randomly assigned. As this is impossible, we have to resort to observational data. To estimate causal effects from observational data, we need to rely on expert knowledge to build the causal graph linking causes, outcome, and other variables [24]. Fig. 1 depicts the graph related to our causal question. We explain our reasoning below.

Our causal question already determines that the causal graph needs to comprise ADAS (the outcome), measures X_1, \dots, X_D (the causes), and the level of beta amyloid 42 peptides ($A\beta$), which determines whether a patient has an Alzheimer’s pathologic change [13]. To link $A\beta$ with the remaining variables, we rely on expert knowledge, namely that $A\beta$ causes levels of Tau phosphorylated at threonine 181 (p-Tau), which in turn causes neurodegeneration that ultimately results in cognitive decline [13,15]. Moreover, we consider that the patient’s levels of $A\beta$ and p-Tau are determined by the allelic variant of apolipoprotein E (ApoE; [21]), among other unobserved common causes (dashed line), and that aging influences the neuroanatomy, amyloid and tau pathology [2,4,26]. Beside




Fig. 1. Causal graph used to estimate the causal effect (red arrow) of subcortical volume (X_d^v) and cortical thickness (X_d^t) on cognitive function (ADAS) in the presence of an unknown and unobserved confounder U . Exogenous variables irrelevant for estimating the causal effect of interest are not shown. Circles are random variables and arrows causal relationships. Filled circles are observed, transparent circles are hidden, bidirectional edges denote unobserved common causes.

biological relationships, we also include demographic and socio-economic factors. In particular, p-Tau levels, brain size, and the level of education is known to differ in males and females [8]. We model resilience to neurodegeneration by including years of education (Edu) as a proxy for cognitive reserve, and total intracranial volume (TIV) as a proxy for brain reserve, where the latter is only causal for volume measurements $X_{D_1}^v, \dots, X_{D_1}^v$ [30]. Finally, fig. 1 reflects that age is a known confounder of the relationship between neuroanatomy and cognition [12]. However, as outlined in the introduction, the full set of confounders is extensive and most of them are unmeasured. Therefore, we assume an unknown and unobserved set of additional confounders U .

2.2 Identifiability in the Presence of an Unobserved Confounder

Formally, the causal question states that we want to use the observed data to estimate the average causal effect that a subset $\mathcal{S} \subset \{X_1, \dots, X_D\}$ of neuroanatomical structures have simultaneously on the ADAS score:

$$\mathbb{E}[\text{ADAS} \mid do(X_{\mathcal{S}} = x'_{\mathcal{S}})] = \int \text{adas} \cdot P(\text{adas} \mid do(x'_{\mathcal{S}})) d\text{adas}. \quad (1)$$

The do-operator states that we are interested in the post-intervention distribution of ADAS, induced by the intervention that sets the neuroanatomical measures $X_{\mathcal{S}}$ equal to $x'_{\mathcal{S}}$. The central question in causal inference is that of identification: Can the post-intervention distribution $P(\text{adas} \mid do(x))$ be estimated from data governed by the observed joint distribution over X and ADAS?

Inspecting fig. 1 reveals that the relationship between any X_d and ADAS is confounded by the known confounder age, but also the unknown confounder U . Knowing that the causal effect of X_d on ADAS is unidentifiable in the presence of unobserved confounding, it initially appears that our causal question cannot be answered [24, Theorem 3.2.5]. Next, we will detail how we can resolve this issue by exploiting the influence that the unknown confounder has on multiple causes simultaneously. The full procedure is outlined in algorithm 1.

Algorithm 1: Causal inference with unobserved confounding.

Input: Neuroanatomical measures $\mathbf{X} \in \mathbb{R}^{N \times D}$, direct influences on neuroanatomy $\mathbf{F} \in \mathbb{R}^{N \times P}$, ADAS scores $\mathbf{y} \in \mathbb{R}^N$; values for intervention x_S ; τ minimum Bayesian p-value for model checking.
Output: Estimate of $\mathbb{E}[\text{ADAS} \mid do(X_S = x_S)]$.

- 1 Sample a random binary matrix $\mathbf{H} \in \{0; 1\}^{N \times D}$ and split data into $\mathbf{X}^{\text{obs}} = (1 - \mathbf{H}) \odot \mathbf{X}$ and $\mathbf{X}^{\text{holdout}} = \mathbf{H} \odot \mathbf{X}$.
- 2 Fit PLFM with parameters θ to \mathbf{X}^{obs} and \mathbf{F} .
- 3 Simulate M replicates of \mathbf{x}_i by drawing from posterior predictive distribution $P(\mathbf{x}_i^{\text{sim}} \mid \mathbf{x}_i^{\text{obs}}) = \int P(\mathbf{x}_i^{\text{sim}} \mid \theta) P(\theta \mid \mathbf{x}_i^{\text{obs}}) d\theta$.
- 4 For each observation, estimate Bayesian p-value using test statistic in (5): $p_{B_i} \approx \frac{1}{M} \sum_{m=1}^M I(T(\mathbf{x}_{i,m}^{\text{sim}}) \geq T(\mathbf{x}_i^{\text{holdout}}))$.
- 5 **if** $\frac{1}{N} \sum_{i=1}^N p_{B_i} > \tau$ **then**
- 6 Estimate substitute confounders $\hat{\mathbf{Z}} = \mathbb{E}[\mathbf{Z} \mid \mathbf{X}^{\text{obs}}, \mathbf{F}]$ by PLFM.
- 7 Fit a regression model $f: \mathbf{x} \mapsto \text{ADAS}$, using the residuals defined in (12).
- 8 $\mathbb{E}[\text{ADAS} \mid do(X_S = x_S)] \approx \frac{1}{N} \sum_{i=1}^N f(r(\tilde{\mathbf{x}}_i, \mathbf{f}_i))$, where $\tilde{\mathbf{x}}_i$ equals \mathbf{x}_i , except for features in \mathcal{S} , which are set to x_S .
- 9 **end**

2.3 Estimating a Substitute Confounder

By assuming that the confounder U is unobserved, we can only attempt to estimate causal effects by building upon assumptions on the data-generating process. Therefore, we assume that the data-generating process is faithful to the graphical model in fig. 1, i.e., statistical independencies in the observed data distribution imply missing causal relationships in the graph. In particular, this implies that there is a common unobserved confounder U that is shared among all causes X and that there is no unobserved confounder that affects a single cause. Here, causes X are image-derived volume and thickness measurements, and we require that the unknown confounder affects multiple brain regions and not just a single region. This assumption is plausible, because common sources of confounding such as scanner, protocol, and aging affect the brain as a whole and not just individual regions [2,31]. Based on this assumption, we can exploit the fact that the confounder induces dependence among multiple causes.

From fig. 1 we can observe that U , age, education, gender, and p-Tau are shared among all causes X_1, \dots, X_D . Given these parents, denoted as PA_{X_1, \dots, X_D} , the causes become conditionally independent:

$$P(x_1, \dots, x_D \mid PA_{X_1, \dots, X_D}) = \prod_{d=1}^D P(x_d \mid PA_{X_1, \dots, X_D}). \tag{2}$$

The key realization of our proposed method is that the conditional probability (2), which is derived solely from the causal graph in fig. 1, has the same form as the conditional distribution of a probabilistic latent factor model (PLFM). Therefore, we can utilize this connection to estimate a substitute confounder \mathbf{z} for the unobserved confounder U via a latent factor model.

The theoretical proof for this approach is due to Wang and Blei [34] who showed that the latent representation of any PLFM does indeed render the relationship between neuroanatomical measures and ADAS unconfounded if (i)

the PLFM captures all multi-cause confounders, and (ii) the PLFM estimates the substitute confounder with consistency, i.e., deterministically, as the number of causes D grows large. To verify whether (i) holds, we rely on posterior predictive checking, as described below, and only proceed with estimation of the causal effect if the check passes.¹ Regarding (ii), [3] showed that estimates of many PLFM are consistent if the number of causes and samples is large. Note, that it does not imply that we need to find the true confounder, just a deterministic bijective transformation of it [34].

Latent Factor Model. Let $\mathbf{f}_i \in \mathbb{R}^P$ be the feature vector describing the observed direct influences on X_d for the i -th patient, except TIV, which we account for by dividing all volume measures $X_1^v, \dots, X_{D_1}^v$ by it. Then, we can use an extended version of probabilistic principal component analysis (PPCA, [32]) to represent the D causes in terms of the known causes \mathbf{f}_i and the latent substitute confounder $\mathbf{z}_i \in \mathbb{R}^K$:

$$\mathbf{x}_i = \mathbf{W}\mathbf{z}_i + \mathbf{A}\mathbf{f}_i + \boldsymbol{\varepsilon}_i, \quad \boldsymbol{\varepsilon}_i \sim \mathcal{N}(\mathbf{0}, \sigma_x^2 \mathbf{I}_D), \quad \forall i = 1, \dots, N, \quad (3)$$

where \mathbf{I}_D is a $D \times D$ identity matrix, \mathbf{W} a $D \times K$ loading matrix, and \mathbf{A} a $D \times P$ matrix of regression coefficients for the known causes of \mathbf{x}_i , excluding TIV.

Our approach is not restricted to PPCA, in fact, any PLFM can be used to infer the substitute confounder. Here, we consider an extended version of probabilistic matrix factorization (BPMF, [27]) as an alternative:

$$x_{ij} = \mathbf{z}_i^\top \mathbf{v}_j + \mathbf{A}_j^\top \mathbf{f}_i + \varepsilon_{ij}, \quad \varepsilon_{ij} \sim \mathcal{N}(0, \sigma_x^2), \quad \forall i = 1, \dots, N, \quad \forall j = 1, \dots, D, \quad (4)$$

where \mathbf{v}_j is a K -dimensional feature-specific latent vector. The full models with prior distributions are depicted in fig. 2.

Posterior predictive checking. To ensure the PLFM can represent the joint distribution over the observed causes well, we employ posterior predictive checking to quantify how well the PLFM fits the data [11, ch. 6]. If the PLFM is a good fit, simulated data generated under the PLFM should look similar to observed data. First, we hold-out a randomly selected portion of the observed causes, yielding \mathbf{X}^{obs} to fit the factor model, and $\mathbf{X}^{\text{holdout}}$ for model checking. Next, we draw simulated data from the joint posterior predictive distribution. If there is a systematic difference between the simulated and the held-out data, we can conclude that the PLFM does not represent the causes well. We use the expected negative log-likelihood as test statistic to compute the Bayesian p-value p_B – the probability that the simulated is more extreme than the observed data [11]:

$$T(\mathbf{x}_i) = \mathbb{E}_\theta [-\log p(\mathbf{x}_i | \theta) | \mathbf{x}_i^{\text{obs}}], \quad p_{B_i} = P(T(\mathbf{x}_i^{\text{sim}}) \geq T(\mathbf{x}_i^{\text{holdout}}) | \mathbf{x}_i^{\text{obs}}), \quad (5)$$

where θ is the set of all parameters of the PLFM. We estimate p_{B_i} by drawing $\mathbf{x}_i^{\text{sim}}$ repeatedly from the posterior predictive distribution and computing the proportion for which $T(\mathbf{x}_i^{\text{sim}}) \geq T(\mathbf{x}_i^{\text{holdout}})$ (see algorithm 1). Next, we will prove that the causal effect of neuroanatomical measures on ADAS is identifiable by accounting for the substitute confounder.

¹ We follow [34] and use $\bar{p}_B = \mathbb{E}[p_{B_i}] > 0.1$ as criterion.

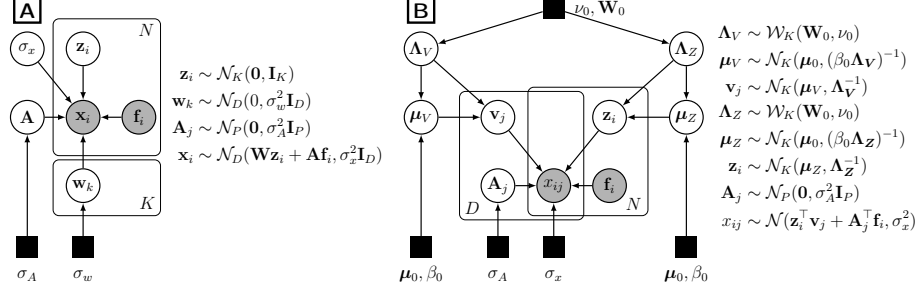


Fig. 2. Probabilistic latent factor models to estimate a K -dimensional substitute confounder \mathbf{z}_i from D -dimensional causes \mathbf{x}_i . Circles are random variables, filled circles are to be observed, transparent circles are to be estimated. A: Probabilistic principal component analysis model. B: Bayesian probabilistic matrix factorization model.

2.4 Identifiability in the Presence of a Substitute Confounder

The theoretical results outlined above and described in detail in [34] enable us to treat the substitute confounder \mathbf{z} as if it were observed. We will now prove that the average causal effect (1) that a subset \mathcal{S} of neuroanatomical structures have simultaneously on the ADAS score is identifiable in this modified setting. Therefore, we will again refer to fig. 1, but replace U with its observed substitute.

Theorem 1. *Assuming the data-generating process is faithful to the graphical model in fig. 1, the causal effect $\mathbb{E}[\text{ADAS} \mid do(x'_\mathcal{S})]$ of a subset \mathcal{S} of neuroanatomical measures on ADAS is identifiable from a distribution over the observed neuroanatomical measures not in \mathcal{S} , age, and the substitute confounder \mathbf{z} .*

Proof. By assuming that the graph in fig. 1 is faithful, we can apply the rules of do calculus [24, Theorem 3.4.1] to show that the post-intervention distribution can be identified from observed data and the substitute confounder. We denote by $\bar{\mathcal{S}}$ the complement of the set \mathcal{S} :

$$\mathbb{E}[\text{ADAS} \mid do(x'_\mathcal{S})] = \mathbb{E}_{age, x_{\bar{\mathcal{S}}}, \mathbf{z}} [\mathbb{E}[\text{ADAS} \mid do(x'_\mathcal{S}), x_{\bar{\mathcal{S}}}, age, \mathbf{z}]] \quad (6)$$

$$= \mathbb{E}_{age, x_{\bar{\mathcal{S}}}, \mathbf{z}} [\mathbb{E}[\text{ADAS} \mid do(x'_\mathcal{S}), x_{\bar{\mathcal{S}}}, do(\text{ptau}), age, \mathbf{z}]] \quad (7)$$

$$= \mathbb{E}_{age, x_{\bar{\mathcal{S}}}, \mathbf{z}} [\mathbb{E}[\text{ADAS} \mid x'_\mathcal{S}, x_{\bar{\mathcal{S}}}, do(\text{ptau}), age, \mathbf{z}]] \quad (8)$$

$$= \mathbb{E}_{age, x_{\bar{\mathcal{S}}}, \mathbf{z}} [\mathbb{E}[\text{ADAS} \mid x'_\mathcal{S}, x_{\bar{\mathcal{S}}}, \text{ptau}, age, \mathbf{z}]] \quad (9)$$

$$= \mathbb{E}_{age, x_{\bar{\mathcal{S}}}, \mathbf{z}} [\mathbb{E}[\text{ADAS} \mid x'_\mathcal{S}, x_{\bar{\mathcal{S}}}, age, \mathbf{z}]] \quad (10)$$

$$\approx \frac{1}{N} \sum_{i=1}^N \hat{\mathbb{E}}[\text{ADAS} \mid x'_\mathcal{S}, \mathbf{x}_{i, \bar{\mathcal{S}}}, age_i, \mathbf{z}_i]. \quad (11)$$

The equality in (6) is due to the factorization given by the graph in fig. 1, the one in (7) due to rule 3 of do calculus, (8) and (9) are due to rule 2 of do calculus, and (10) is due to $\text{ADAS} \perp\!\!\!\perp \text{p-Tau} \mid x'_\mathcal{S}, x_{\bar{\mathcal{S}}}, age, \mathbf{z}$. Finally, we can estimate the outer expectation by Monte Carlo and the inner expectation with a regression model from the observed data alone, if $P(x_\mathcal{S} \mid PA_{X_1, \dots, X_D}) > 0$ for any subset \mathcal{S} . This assumption holds for the proposed PLFM in (3) and (4), because their conditional distribution is a normal distribution, which is non-zero everywhere.

Table 1. RMSE $\times 100$ of effects estimated by logistic regression compared to the true causal effects on semi-synthetic data. ROA is the error when only regressing out the observed confounder age, Oracle is the error when including all confounders. Columns with Δ denote the improvement over the ‘Non-causal’ model.

	ν_x/ν_z	Non-causal	ROA	PPCA	BPMF	Oracle	Δ ROA	Δ PPCA	Δ BPMF
least confounded	10/1	20.121	19.416	17.449	18.033	17.917	0.705	2.672	2.087
	5/1	21.078	20.436	18.560	18.895	18.781	0.643	2.518	2.183
	4/1	21.505	20.889	19.050	19.296	19.169	0.617	2.456	2.210
	3/1	22.169	21.590	19.820	19.933	19.782	0.580	2.349	2.236
	5/2	22.653	22.097	20.382	20.408	20.233	0.556	2.270	2.244
	5/3	23.911	23.417	21.837	21.699	21.438	0.494	2.073	2.212
	3/2	24.275	23.798	22.261	22.083	21.796	0.477	2.014	2.192
	1/1	25.802	25.391	24.022	23.735	23.319	0.411	1.780	2.067
	2/3	27.464	27.116	25.932	25.596	25.040	0.348	1.531	1.867
	3/5	27.899	27.567	26.434	26.088	25.502	0.333	1.465	1.811
most confounded	2/5	29.564	29.284	28.354	28.018	27.307	0.280	1.210	1.546
	1/3	30.297	30.037	29.199	28.877	28.124	0.259	1.097	1.419
	1/4	31.384	31.153	30.448	30.167	29.354	0.230	0.936	1.216
	1/5	32.179	31.967	31.364	31.118	30.266	0.212	0.815	1.060
	1/10	34.318	34.154	33.816	33.699	32.778	0.164	0.502	0.620

2.5 The Outcome Model

Theorem 1 tells us that the average causal effect can be estimated from the observed data and the substitute confounder. The final step in causal inference is to actually estimate the expectation in (11), which we do by fitting a model to predict the ADAS score from neuroanatomical measures, age and the substitute confounder (see algorithm 1). We will use a linear model, but Theorem 1 holds when using a non-linear model too.

ADAS ranges between 0 and 85 with higher values indicating a higher cognitive decline, hence we convert ADAS to proportions in the interval $(0, 1)$ and use a Bayesian Beta regression model for prediction [7]. Let y_i denote the ADAS score of subject i , then the likelihood function is

$$L(\beta_0, \boldsymbol{\beta}, \phi) = \prod_{i=1}^N \frac{y_i^{\mu_i \phi - 1} (1 - y_i)^{(1 - \mu_i) \phi - 1}}{B(\mu_i \phi, (1 - \mu_i) \phi)}, \quad \mu_i = \text{logit}^{-1}(\beta_0 + r(\mathbf{x}_i, \mathbf{f}_i)^\top \boldsymbol{\beta}),$$

where B is the beta function, and ϕ is a scalar scale parameter. To account for unobserved confounding, we replace the original neuroanatomical measures by the residuals with respect to their reconstruction by the PLFM:

$$r(\mathbf{x}_i, \mathbf{f}_i) = \mathbf{x}_i - \mathbb{E}[X_1, \dots, X_D \mid \hat{\mathbf{z}}_i, \mathbf{f}_i], \quad \hat{\mathbf{z}}_i = \mathbb{E}[Z \mid \mathbf{x}_i^{\text{obs}}, \mathbf{f}_i], \quad (12)$$

where the expectations are with respect to the selected PLFM.²

3 Experiments

Semi-synthetic Data. In our first experiment, we evaluate how well causal effects can be recovered when ignoring all confounders, using only the observed

² Code available at <https://github.com/ai-med/causal-effects-in-alzheimers-continuum>

confounder (age), using the observed and unobserved confounders (oracle), and using the observed and substitute confounder computed by (3) or (4). We use T1-weighted magnetic resonance imaging brains scans from $N = 11,800$ subjects from UK Biobank [22]. From each scan, we extract 19 volume measurements with FreeSurfer 5.3 [9] and create a synthetic binary outcome. For each volume, we use age and gender as known causes \mathbf{f}_i and estimate $\hat{x}_{ij} = \mathbb{E}[X_{ij} | \mathbf{f}_i]$ via linear regression. We use age as an observed confounder and generate one unobserved confounder u_k by assigning individuals to clusters with varying percentages of positive labels. First, we obtain the first two principal components across all volumes, scale individual scores to $[0; 1]$, and cluster the projected data into 4 clusters using k -means to assign $u_k \in \{1, 2, 3, 4\}$ and scale σ_k of the noise term. Causal effects follow a sparse normal distribution (\mathcal{N}_{sp}), where all values in the 20–80th percentile range are zero, hence only a small portion of volumes have a non-zero causal effect. Let $\nu_x, \nu_z, \nu_\varepsilon = 1 - \nu_x - \nu_z$ denote how much variance can be explained by the causal effects, confounding effects, and the noise, respectively, then for the i -th instance in cluster k , we generate a binary outcome y_i as:

$$\begin{aligned}
 y_i &\sim \text{Bernoulli}(\text{logit}^{-1}(\beta_0 + \hat{\mathbf{x}}_i \boldsymbol{\beta} \frac{\sqrt{\nu_x}}{\sigma_x} + u_k \cdot \frac{\sqrt{0.9\nu_z}}{\sigma_z} + \text{age}_i \gamma \cdot \frac{\sqrt{0.1\nu_z}}{\sigma_{\text{age}}} + \varepsilon_i \frac{\sqrt{\nu_\varepsilon}}{\sigma_\varepsilon})), \\
 \beta_j &\sim \mathcal{N}_{\text{sp}}(0, 0.5), \quad \gamma \sim \mathcal{N}(0, 0.2), \quad \sigma_k \sim 1 + \text{InvGamma}(3, 1), \quad \varepsilon_i \sim \mathcal{N}(0, \sigma_k),
 \end{aligned}$$

where $\sigma_x, \sigma_z, \sigma_{\text{age}}$, and σ_ε are standard deviations with respect to $\hat{\mathbf{x}}_i \boldsymbol{\beta}, u, \text{age}_i$ and ε_i for $i = 1, \dots, N$. Finally, we choose β_0 such that the positive and negative class are roughly balanced.

Table 1 shows the root mean squared error (RMSE) with respect to the true causal effect of a logistic regression model across 1,000 simulations for various ν_x, ν_z , and $\nu_\varepsilon = 0.1$. We used $K = 5$ substitute confounders and both PLFM passed the posterior predictive check with $\bar{p}_B = 0.269$ (BPMF) and $\bar{p}_B = 0.777$ (PPCA), despite that the true data generation model differs. As expected, by ignoring confounding completely (first column), the RMSE is the highest. When only accounting for the known confounder age (second column), the RMSE decreases slightly. The RMSE reduces considerably when using a substitute confounder and achieves an improvement 2.9 – 5.5 higher than that of the age-only model. Finally, the results show that there is a cost to using a substitute confounder: using all confounders (Oracle) leads to the lowest RMSE.

Alzheimer’s Disease Data. In this experiment, we study the causal effect of neuroanatomical measures on ADAS using data from the Alzheimer’s Disease Neuroimaging Initiative [14]. We only focus on effects due to Alzheimer’s pathologic change and not other forms of dementia. Therefore, we only include patients with abnormal amyloid biomarkers [13]. We extract 14 volume and 8 thickness measures using FreeSurfer [9] for 711 subjects (highly correlated measures are removed). Since the average causal effect is fully parameterized by the coefficients of the linear Beta regression model, we can compare estimated coefficients of the proposed approach with $K = 6$ substitute confounders, with that of a model ignoring all confounders (Non-causal), and of a model trained on measures where age, gender, and education has been regressed-out.

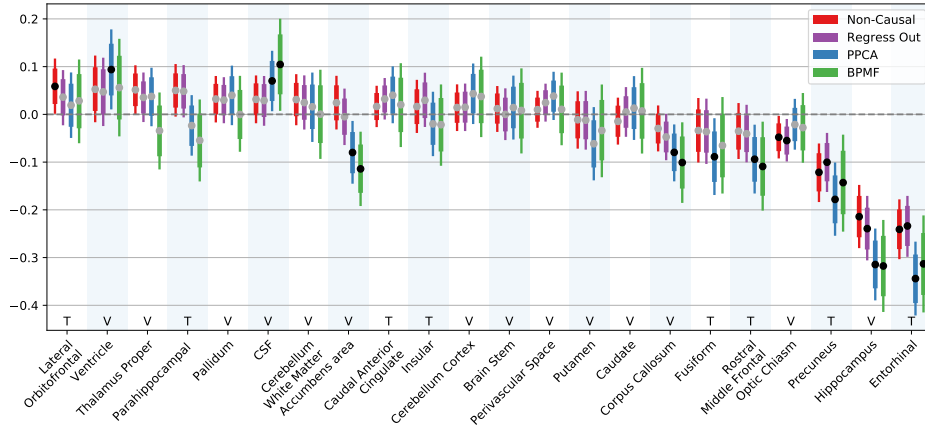


Fig. 3. Mean coefficient (dot), 80% (thick line), and 95% (thin line) credible interval of volume and thickness measures. Significant effects are marked with a black dot.

The BPFM ($\bar{p}_B = 0.293$) and PPCA model ($\bar{p}_B = 0.762$) passed the posterior predictive check; estimated coefficients are depicted in fig. 3. Lateral orbitofrontal thickness and optic chiasm become non-significant after correcting for unobserved confounding, whereas rostral anterior cingulate thickness, CSF volume, accumbens volume, and corpus callosum volume become significant. The biggest change concerns accumbens volume, which is associated with cognitive *improvement* in the non-causal model, but is a cause for cognitive *decline* in the causal models. This correction is justified, because the accumbens is part of the limbic circuit and thus shares its vulnerability to degenerate during cognitive decline [16]. An analysis based on the non-causal model would have resulted in a wrong conclusion. The result that atrophy of corpus callosum is causal seems to be plausible too, because it is a known marker of the progressive interhemispheric disconnection in AD [5]. While finding literature on the absence of an effect on cognition is difficult, we believe a causal effect of atrophy of the optic nerve to be unlikely, because its main function is to transmit visual information. It is more likely that the non-causal model picked up an association due to aging-related confounding instead. Finally, we want to highlight the change in parahippocampal thickness. Previous research suggests that thinning is common in AD [28], which would only be captured correctly after correcting for unobserved confounding, as indicated by a negative mean coefficient. In contrast, when only accounting for the known confounder age, the estimated mean coefficient remains positive.

4 Conclusion

Inferring causal effects from observational neuroimaging data is challenging, because it requires expressing the causal question in a graphical model, and various unknown sources of confounding often render the causal effects unidentifiable. We tackled this task by deriving a causal graph from the current clinical knowledge on

the Alzheimer’s disease continuum, and proposed a latent factor model approach to estimate a substitute for the unobserved confounders. Our experiments on semi-synthetic data showed that our proposed approach can recover causal effects more accurately than a model that only accounts for observed confounders or ignores confounding. Analyses on the causal effects on cognitive decline due to Alzheimer’s revealed that accounting for unobserved confounding reveals important causes of cognitive decline that otherwise would have been missed.

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, and the Federal Ministry of Education and Research in the call for Computational Life Sciences (DeepMentia, 031L0200A).

References

1. Alfaró-Almagro, F., McCarthy, P., Afyouni, S., Andersson, J.L.R., Bastiani, M., et al.: Confound modelling in UK Biobank brain imaging. *Neuroimage* **224**, 117002 (2021)
2. Barnes, J., Ridgway, G.R., Bartlett, J., Henley, S.M., Lehmann, M., et al.: Head size, age and gender adjustment in MRI studies: a necessary nuisance? *Neuroimage* **53**(4), 1244–1255 (2010)
3. Chen, Y., Li, X., Zhang, S.: Structured latent factor analysis for large-scale data: Identifiability, estimability, and their implications. *J Am Stat Assoc* pp. 1–15 (2019)
4. Crary, J.F., Trojanowski, J.Q., Schneider, J.A., Abisambra, J.F., Abner, E.L., et al.: Primary age-related tauopathy (PART): a common pathology associated with human aging. *Acta Neuropathol* **128**(6), 755–766 (2014)
5. Delbeuck, X., Van der Linden, M., Collette, F.: Alzheimer’s Disease as a Disconnection Syndrome? *Neuropsychol Rev* **13**(2), 79–92 (2003)
6. Dukart, J., Schroeter, M.L., Mueller, K.: Age correction in dementia—matching to a healthy brain. *PloS One* **6**(7), e22193 (2011)
7. Ferrari, S., Cribari-Neto, F.: Beta regression for modelling rates and proportions. *J Appl Stat* **31**(7), 799–815 (2004)
8. Ferretti, M.T., Martinkova, J., Biskup, E., Benke, T., Gialdini, G., et al.: Sex and gender differences in Alzheimer’s disease: current challenges and implications for clinical practice. *Eur J Neurol* **27**(6), 928–943 (2020)
9. Fischl, B.: FreeSurfer. *Neuroimage* **62**(2), 774–781 (2012)
10. Fortin, J.P., Cullen, N., et al.: Harmonization of cortical thickness measurements across scanners and sites. *Neuroimage* **167**, 104–120 (2018)
11. Gelman, A., Carlin, J.B., Stern, H.S., Rubin, D.B., Dunson, D.B.: *Bayesian Data Analysis*. Taylor & Francis, 3rd edn. (2014)
12. Hedden, T., Gabrieli, J.D.E.: Insights into the ageing mind: a view from cognitive neuroscience. *Nat Rev Neurosci* **5**(2), 87–96 (2004)
13. Jack, C.R., Bennett, D.A., Blennow, K., Carrillo, M.C., Dunn, B., et al.: NIA-AA Research Framework: Toward a biological definition of Alzheimer’s disease. *Alzheimers Dement* **14**(4), 535–562 (2018)
14. Jack, C.R., Bernstein, M.A., Fox, N.C., Thompson, P., et al.: The Alzheimer’s disease neuroimaging initiative (ADNI): MRI methods. *J Magn Reson Imaging* **27**(4), 685–691 (2008)

15. Jack, C.R., Knopman, D.S., Jagust, W.J., Petersen, R.C., Weiner, M.W., et al.: Tracking pathophysiological processes in Alzheimer’s disease: an updated hypothetical model of dynamic biomarkers. *Lancet Neurol* **12**(2), 207–216 (2013)
16. de Jong, L.W., Wang, Y., White, L.R., Yu, B., van Buchem, M.A., Launer, L.J.: Ventral striatal volume is associated with cognitive decline in older people: a population based MR-study. *Neurobiol Aging* **33**, 424.e1–424.10 (2012)
17. Koikkalainen, J., Pölönen, H., Mattila, J., van Gils, M., Soininen, H., Lötjönen, J., et al.: Improved classification of Alzheimer’s disease data via removal of nuisance variability. *PLoS One* **7**(2), e31112 (2012)
18. Kostro, D., Abdulkadir, A., Durr, A., Roos, R., Leavitt, B.R., et al.: Correction of inter-scanner and within-subject variance in structural MRI based automated diagnosing. *Neuroimage* **98**, 405–415 (2014)
19. Linn, K.A., Gaonkar, B., Doshi, J., Davatzikos, C., Shinohara, R.T.: Addressing confounding in predictive models with an application to neuroimaging. *Int J Biostat* **12**(1), 31–44 (2016)
20. Lockhart, S.N., DeCarli, C.: Structural imaging measures of brain aging. *Neuropsychol Rev* **24**(3), 271–289 (2014)
21. Long, J.M., Holtzman, D.M.: Alzheimer Disease: An Update on Pathobiology and Treatment Strategies. *Cell* **179**(2), 312–339 (2019)
22. Miller, K.L., Alfaro-Almagro, F., Bangerter, N.K., Thomas, D.L., Yacoub, E., et al.: Multimodal population brain imaging in the UK Biobank prospective epidemiological study. *Nat Neurosci* **19**(11), 1523–1536 (2016)
23. Mohs, R.C., Knopman, D., Petersen, R.C., Ferris, S.H., Ernesto, C., et al.: Development of cognitive instruments for use in clinical trials of antidementia drugs. *Alzheimer Dis Assoc Disord* **11**, 13–21 (1997)
24. Pearl, J.: *Causality: Models, reasoning, and inference*. Cambridge University Press (2000)
25. Rao, A., Monteiro, J.M., Mourao-Miranda, J.: Predictive modelling using neuroimaging data in the presence of confounds. *Neuroimage* **150**, 23–49 (2017)
26. Rodrigue, K.M., Kennedy, K.M., Park, D.C.: Beta-Amyloid Deposition and the Aging Brain. *Neuropsychol Rev* **19**(4), 436–450 (2009)
27. Salakhutdinov, R., Mnih, A.: Bayesian Probabilistic Matrix Factorization Using Markov Chain Monte Carlo. In: *ICML*. pp. 880–887 (2008)
28. Schwarz, C.G., Gunter, J.L., Wiste, H.J., Przybelski, S.A., Weigand, S.D., et al.: A large-scale comparison of cortical thickness and volume methods for measuring alzheimer’s disease severity. *Neuroimage Clin* **11**, 802–812 (2016)
29. Snoek, L., Miletic, S., Scholte, H.S.: How to control for confounds in decoding analyses of neuroimaging data. *Neuroimage* **184**, 741–760 (2019)
30. Stern, Y., Arenaza-Urquijo, E.M., Bartrés-Faz, D., Belleville, S., Cantilon, M., et al.: Whitepaper: Defining and investigating cognitive reserve, brain reserve, and brain maintenance. *Alzheimer’s Dement* **16**, 1305–1311 (2020)
31. Stonnington, C.M., Tan, G., Klöppel, S., Chu, C., Draganski, B., Jack, C.R., et al.: Interpreting scan data acquired from multiple scanners: A study with Alzheimer’s disease. *Neuroimage* **39**(3), 1180–1185 (2008)
32. Tipping, M.E., Bishop, C.M.: Probabilistic Principal Component Analysis. *J R Stat Soc Series B Stat Methodol* **61**(3), 611–622 (1999)
33. Wachinger, C., Rieckmann, A., Pölsterl, S.: Detect and Correct Bias in Multi-Site Neuroimaging Datasets. *Med Image Anal* (2020)
34. Wang, Y., Blei, D.M.: The Blessings of Multiple Causes. *J Am Stat Assoc* pp. 1–71 (2019)