

Federated Causal Inference on Multi-Site Observational Data via Propensity Score Aggregation

Rémi Khellaf¹ Aurélien Bellet¹ Julie Josse¹

Abstract

Causal inference typically assumes centralized access to individual-level data. Yet, in practice, data are often decentralized across multiple sites, making centralization infeasible due to privacy, logistical, or legal constraints. We address this problem by estimating the Average Treatment Effect (ATE) from decentralized observational data via a Federated Learning (FL) approach, allowing inference through the exchange of aggregate statistics rather than individual-level data. We propose a novel method to estimate propensity scores via a federated weighted average of local scores using Membership Weights (MW), defined as probabilities of site membership conditional on covariates. MW can be flexibly estimated with parametric or non-parametric classification models using standard FL algorithms. The resulting propensity scores are used to construct Federated Inverse Propensity Weighting (Fed-IPW) and Augmented IPW (Fed-AIPW) estimators. In contrast to meta-analysis methods, which fail when any site violates positivity, our approach exploits heterogeneity in treatment assignment across sites to improve overlap. We show that Fed-IPW and Fed-AIPW perform well under site-level heterogeneity in sample sizes, treatment mechanisms, and covariate distributions. Theoretical analysis and experiments on simulated and real-world data demonstrate clear advantages over meta-analysis and related approaches.

1. Introduction

The Average Treatment Effect (ATE) is a key causal estimand used to quantify the effect of a treatment on an outcome and is commonly employed as the primary measure of efficacy in evaluating new therapies, including vaccines, before regulatory approval (Polack et al., 2020). In Randomized Clinical Trials (RCTs), treatment assignment is randomized, ensuring that the observed association between treatment and outcome reflects a causal effect. Under this design, the ATE can be consistently estimated using a simple Difference-in-Means (DM) estimator (Splawa-Neyman, 1990), which can be further refined through covariate adjustment to reduce variance (FDA, 2023; EMA, 2024; Lei & Ding, 2021). However, RCTs are often expensive, time-consuming, or infeasible. In such cases, estimating treatment effects from observational data becomes the only viable alternative (Hernán, 2018; Hernán & Robins, 2006). Although such real-world data is abundant, drawing causal inferences from it is challenging due to confounding covariates, rendering the unadjusted DM estimator biased (Grimes & Schulz, 2002). Adjusting for confounders is thus essential (VanderWeele, 2019). This can be done by predicting counterfactual outcomes before averaging the differences (the G-formula plug-in estimator, Robins, 1986). Another approach is to weight individuals according to their treatment probability, emulating a randomized trial. For instance, the Inverse Propensity Weighting (IPW) estimator (Rosenbaum & Rubin, 1983) relies on estimating the *propensity score*—the probability of treatment given covariates. Doubly robust estimators such as the Augmented IPW (AIPW) (Bang & Robins, 2005) combine weighting with outcome modeling to remain consistent as long as either model is correctly specified.

Larger datasets improve the precision of treatment effect estimates, especially for underrepresented subgroups. Yet real-world data is typically decentralized—spread across hospitals, companies, or countries—making aggregation difficult, particularly in healthcare where privacy regulations, data ownership, and governance issues impede centralization. Federated Learning (FL) (Kairouz et al., 2021) offers a solution to train models across distributed data without sharing individual-level data. Although FL has

¹PreMeDiCaL team, Inria, Inserm, Idesp, Université de Montpellier, France. Correspondence to: Rémi Khellaf <remi.khellaf@inria.fr>.

been widely applied to prediction tasks, its extension to causal inference is more recent and has emerged as a rapidly growing area of research. This problem is especially challenging in observational studies, where differences in covariate distributions and treatment assignment mechanisms across sites create multiple sources of heterogeneity that must be addressed without sharing raw data, while achieving results comparable to centralized analyses—an issue that remains largely unsolved.

Contributions. We propose federated (A)IPW estimators for decentralized observational data, moving beyond the aggregation of local ATE estimates used in meta-analysis (Riley et al., 2023). At the core of our approach is a flexible, primarily non-parametric strategy for federating propensity scores. Unlike prior methods that fit a single global parametric model via parameter averaging (Xiong et al., 2023) or federated gradient descent (Guo et al., 2025), we construct a global propensity score as a mixture of locally estimated models. This involves (1) local estimation of propensity scores at each site—accommodating heterogeneity in treatment assignment and flexibility in model choice—and (2) aggregation into a global score using Membership Weights (MW), i.e., the probability of site membership given covariates. MW can be estimated in a federated manner using flexible, potentially non-parametric classification models, ensuring both robustness and communication efficiency. In contrast, transferring Density Ratio Weights (DW) from the transportability and generalization literature to our setting requires strong modeling assumptions, as they are otherwise incompatible with FL constraints. Using our federated propensity scores, we build the Federated IPW (Fed-IPW) and its augmented variant (Fed-AIPW), derive their variances, and show they achieve equal or lower variance than meta-analysis estimators.

Our approach is particularly advantageous when overlap between treatment groups is poor or absent within sites. In such scenarios, cross-site collaboration becomes crucial, as combining sites increases overall overlap and enables treatment effect estimation that may be infeasible locally. Indeed, when treatment assignment mechanisms differ substantially—such as when one site treats a subgroup absent elsewhere—the combined dataset achieves markedly greater overlap, allowing treatment effects to be estimated that would otherwise be poorly identified in isolation. Additionally, our framework naturally accommodates heterogeneity in sample sizes, treatment policies, covariate distributions, and violations of positivity. Numerical experiments on simulated and real data confirm our theoretical findings and highlight the method’s practical benefits.

Related work. Federated causal inference is still nascent. Khellaf et al. (2025) estimate federated G-formula ATE estimates across multiple RCTs by fitting parametric out-

come models at each site via FL, achieving lower variance than DM. In observational settings, Vo et al. (2022) propose a federated Bayesian approach using Gaussian processes with a shared covariance kernel, but it requires sharing the first four moments of the data, limiting scalability and efficiency. Guo et al. (2025) learn a global propensity score via consensus voting over parametric parameters, retaining only sites meeting a shared specification. In contrast, our method assumes no common propensity score: each site may fit its own model. To address site heterogeneity, Xiong et al. (2023) use a logistic propensity model with shared and site-specific parameters, federating only the common ones. Yin et al. (2025) fit a global model adjusting for covariates and site membership but limit heterogeneity to a site-specific scalar. By contrast, our method makes no structural assumptions, enabling fully nonparametric estimation with heterogeneous local models and relaxing the need for local overlap at each site.

A related line of work focuses on generalizing causal findings from multi-site source populations to a target population. Han et al. (2025) use density ratio weighting of local ATEs to adjust for covariate shift but assume homogeneous nuisance functions across sites and rely on meta-analysis of aggregate statistics. Guo et al. (2024) extend this approach by applying density-ratio weights to combine local propensity scores into a target-specific score. While their framework can target the super-population of the K sites as a special case, it is significantly more constrained than our approach: it requires access to individual-level target covariates via an external public dataset, reweights each treatment arm separately within each site (needing large per-arm sample sizes), and relies on direct estimation of density ratios either through strongly parametric logistic models—which are difficult to validate and sensitive to misspecification—or via nonparametric approaches (e.g., kernel-based methods) that involve sharing raw covariates, making them incompatible with federated learning.

To summarize, our approach stands out for its flexibility and ease of implementation in a federated context: (i) It allows local propensity scores to take arbitrary forms and vary freely across sites, unlike Xiong et al. (2023); Han et al. (2025), which rely on strong parametric assumptions, sometimes with shared parameters; (ii) It truly satisfies the federated learning requirement that no raw data be shared, unlike density ratio-based methods that are either highly parametric or lack federated algorithms (Yin et al., 2025; Han et al., 2025; Guo et al., 2024); and (iii) It does not require local overlap, unlike (Xiong et al., 2023; Han et al., 2025), and to the best of our knowledge, is the first to provably improve overlap relative to meta-analysis.

2. Preliminaries

2.1. ATE Estimators from Centralized Multi-Site Observational Data

In this section, we recall the key components of ATE estimation in a centralized multi-site setting. Following the potential outcomes framework (Rubin, 1974; Splawa-Neyman, 1990), we consider random variables $(X, H, W, Y(1), Y(0))$, where $X \in \mathbb{R}^d$ represents patient covariates, $H \in [K]$ indicates site membership, $W \in \{0, 1\}$ denotes the binary treatment, and $Y(1)$ and $Y(0)$ are the potential outcomes under treatment and control, respectively. We assume that the Stable Unit Treatment Values Assumption (SUTVA) holds, so that the observed outcome is $Y = WY(1) + (1 - W)Y(0)$, and that the potential outcomes are uniformly bounded. In the centralized setting, we observe $n = \sum_{k=1}^K n_k$ observations of independently and identically distributed (i.i.d.) tuples $(H_i, X_i, W_i, Y_i)_{i \in 1, \dots, n} \sim \mathcal{P}^{\otimes n}$, with $n_k = \sum_{i=1}^n \mathbb{1}_{\{H_i=k\}}$ the number of observations in site k . We aim to estimate the ATE defined as the risk difference $\tau = \mathbb{E}[Y(1) - Y(0)] = \mathbb{E}[\mathbb{E}[Y(1) - Y(0) | H]]$, where the expectation is taken over the population $\mathcal{P}$. To be able to identify the ATE, we assume unconfoundedness (standard in causal inference) and further consider Assumption 2, which is specific to the multi-site setting. We provide a federated test for this assumption in Section G.1.

Assumption 1 (Unconfoundedness).

$$\{Y(0), Y(1)\} \perp\!\!\!\perp W \mid (X, H).$$

Assumption 2 (Ignorability on sites).

$$\{Y(0), Y(1)\} \perp\!\!\!\perp H \mid X.$$

Remark (Moment requirements). While Assumption 2 states distributional exchangeability across sites—requiring $Y(w)|X$ to be invariant across sites, consistent ATE estimation only requires *mean*-exchangeability, i.e., $\forall w \in \{0, 1\}$, $\mathbb{E}\{Y(w) \mid X, H\} = \mathbb{E}\{Y(w) \mid X\}$. Variance calculations further require equality of the second moments across sites $\mathbb{E}\{Y^2(w) \mid X, H\} = \mathbb{E}\{Y^2(w) \mid X\}$, i.e., homoscedasticity. We study relaxations of this condition under heteroscedastic site effects in Section B.

Assumptions 1 and 2 imply $\{Y(0), Y(1)\} \perp\!\!\!\perp H \mid (X, W)$ (Dawid, 1979, Lemma 4.3). This entails the *no center-outcome association* condition (Robertson et al., 2021), yielding the testable null hypothesis $\mathbb{E}[Y \mid X, W, H] = \mathbb{E}[Y \mid X, W]$. This ensures that X forms a sufficient set of covariates for confounding adjustment. Our setting is depicted in the graphical model in Figure 1, highlighting that we exclude any direct effect of the site on the outcome. In Section 3.3, we introduce a relaxation of Assumption 2 and extend our approach to handle the presence of site effects.




Figure 1. Graphical model for multi-site observational data.

We define $\mu_w(x) = \mathbb{E}[Y \mid X = x, W = w]$ for $w \in \{0, 1\}$, and let $\tau(x) = \mu_1(x) - \mu_0(x) = \mathbb{E}[Y(1) - Y(0) \mid X = x]$ be the Conditional Average Treatment Effect (CATE). The oracle propensity score is denoted by $e(x) = \mathbb{P}(W \mid X = x)$, and we consider the weak (global) overlap assumption (Wager, 2024).

Assumption 3 (Global overlap). $\mathcal{O}_{\text{global}} = \mathbb{E}[(e(X)(1 - e(X)))^{-1}] < +\infty$.

Assumption 3 is crucial for propensity score-based estimators, as it states that every region of the covariate space has a non-zero probability of receiving both treatments. A lower value of $\mathcal{O}_{\text{global}}$ indicates that these probabilities lie further away from 0 and 1, which corresponds to better overlap. For further insights on overlap, see Li et al. (2018a;b), and for a “misoverlap” metric, refer to Clivio et al. (2024).

With Assumptions 1, 2 and 3, the ATE is identifiable as $\tau = \mathbb{E}\left[\frac{WY}{e(X)} - \frac{(1-W)Y}{1-e(X)}\right]$ (see Appendix A.1). Throughout the paper, we denote oracle ATE estimators, which assume knowledge of the nuisance components e, μ_0, μ_1 , by a superscript *. We define the *Oracle multi-site centralized estimators* as follows:

$$\hat{\tau}_{\text{IPW}}^* = \frac{1}{n} \sum_{i=1}^n \tau_{\text{IPW}}(X_i; e), \quad (1)$$

$$\hat{\tau}_{\text{AIPW}}^* = \frac{1}{n} \sum_{i=1}^n \tau_{\text{AIPW}}(X_i; e, \mu_1, \mu_0), \quad (2)$$

where $\tau_{\text{IPW}}(X_i; e) = \frac{W_i Y_i}{e(X_i)} - \frac{(1-W_i)Y_i}{1-e(X_i)}$ and

$$\begin{aligned} \tau_{\text{AIPW}}(X_i; e, \mu_1, \mu_0) &= \mu_1(X_i) - \mu_0(X_i) \\ &+ \frac{W_i(Y_i - \mu_1(X_i))}{e(X_i)} - \frac{(1-W_i)(Y_i - \mu_0(X_i))}{1-e(X_i)}. \end{aligned}$$

These estimators are unbiased and asymptotically normal.

Theorem 1. Under Assumptions 1, 2 and 3, we have $\sqrt{n}(\hat{\tau} - \tau) \rightarrow \mathcal{N}(0, V)$ with

$$\begin{cases} V_{\text{IPW}} = \mathbb{E}\left[\frac{Y(1)^2}{e(X)}\right] + \mathbb{E}\left[\frac{Y(0)^2}{1-e(X)}\right] - \tau^2, \\ V_{\text{AIPW}} = \mathbb{V}[\tau(X)] \\ \quad + \mathbb{E}\left[\left(\frac{Y - \mu_1(X)}{e(X)}\right)^2\right] + \mathbb{E}\left[\left(\frac{Y - \mu_0(X)}{1-e(X)}\right)^2\right]. \end{cases}$$

The above asymptotic variances align with the single-site setting (Hirano et al., 2003), as detailed in Appendix A.2. However, in practice, the propensity score and outcome models are typically unknown and must be estimated from data. This creates a challenge in the decentralized setting, where centralizing data to compute μ_1 , μ_0 , and e is not feasible. Therefore, the estimators in Definition 1 need to be adapted to this setting. Importantly, the (non-oracle) AIPW estimator is inherently *doubly robust*, remaining consistent as long as either the outcome or the propensity score model is correctly specified (Chernozhukov et al., 2018).

2.2. Meta-Analysis Estimators

We now turn to a decentralized setting in which the K sites cannot share individual-level data. A natural baseline for estimating the ATE across sites is a two-stage meta-analysis approach (Burke et al., 2017), wherein each site independently estimates the relevant nuisance parameters and communicates only the resulting ATE estimates for aggregation. In this setting, we need the following assumption.

Assumption 4 (Local overlap).

$$\forall k \in [K], \mathcal{O}_k = \mathbb{E}[(e(X)(1-e(X)))^{-1} | H = k] < +\infty.$$

Assumption 4 is much stronger than global overlap (Assumption 3), as it must hold at every site. Denoting by $e_k(x) = \mathbb{P}(W = 1 | X = x, H = k)$ the oracle local propensity score at site k , we can define the oracle meta-analysis estimators as follows:

$$\hat{\tau}_{\text{IPW}}^{\text{meta}*} = \sum_{k=1}^K \frac{n_k}{n} \hat{\tau}_{\text{IPW}}^{(k)}, \quad \hat{\tau}_{\text{AIPW}}^{\text{meta}*} = \sum_{k=1}^K \frac{n_k}{n} \hat{\tau}_{\text{AIPW}}^{(k)}, \quad (3)$$

where $\hat{\tau}_{\text{IPW}}^{(k)} = \frac{1}{n_k} \sum_{i=1}^{n_k} \tau_{\text{IPW}}(X_i; e_k)$ and $\hat{\tau}_{\text{AIPW}}^{(k)} = \frac{1}{n_k} \sum_{i=1}^{n_k} \tau_{\text{AIPW}}(X_i; e_k, \mu_1, \mu_0)$ are the local estimators at site k . While alternative aggregation weights—such as the inverse variance of local estimates—can be considered, they produce, in our setting, biased estimates of the global ATE $\tau = \sum_{k=1}^K \rho_k \tau_k$, where $\rho_k = \mathbb{P}(H = k)$ and $\tau_k = \mathbb{E}[Y(1) - Y(0) | H = k]$ is the local ATE. This bias appears whenever the τ_k differ, which commonly occurs when covariate distributions vary across sites and treatment effects are heterogeneous (i.e., depend on covariates), see (Berenfeld et al., 2025).

Theorem 2. *Under Assumptions 1, 2 and 4, the oracle meta-analysis estimators are unbiased for the ATE with asymptotic variances*

$$V_{\text{IPW}}^{\text{meta}*} = \sum_{k=1}^K \rho_k V_{\text{IPW}}^{(k)} + \mathbb{V}[\tau_H],$$

$$V_{\text{AIPW}}^{\text{meta}*} = \sum_{k=1}^K \rho_k V_{\text{AIPW}}^{(k)} + \mathbb{V}[\tau_H],$$

with within-site variance

$$\begin{cases} V_{\text{IPW}}^{(k)} = \mathbb{E} \left[\frac{Y(1)^2}{e_k(X)} | H = k \right] + \mathbb{E} \left[\frac{Y(0)^2}{1-e_k(X)} | H = k \right] - \tau_k^2 \\ V_{\text{AIPW}}^{(k)} = \mathbb{V}[\tau(X) | H = k] + \mathbb{E} \left[\left(\frac{Y - \mu_1(X)}{e_k(X)} \right)^2 | H = k \right] \\ \quad + \mathbb{E} \left[\left(\frac{Y - \mu_0(X)}{1-e_k(X)} \right)^2 | H = k \right], \end{cases}$$

and $\mathbb{V}[\tau_H] = \mathbb{V}[\mathbb{E}[Y(1) - Y(0) | H]]$ the between-sites variance of the local ATEs.

This result is proved in Appendix A.3. A key limitation of meta-analysis estimators is their reliance on Assumption 4, which is fragile and often violated—for instance, when a site applies a deterministic treatment policy (treating all patients or only a subgroup). In such cases, these estimators are ill-defined, yielding biased ATE estimates. To address this, we propose a federated approach that constructs the global propensity score e as a weighted combination of local scores e_k , enabling valid inference even without local overlap.

3. Federated Estimators via Propensity Score Aggregation

3.1. Oracle Federated Estimators

As discussed before, existing federated causal inference methods often rely on restrictive assumptions—such as a common propensity score across sites (Guo et al., 2025), site differences limited to intercept shifts (Yin et al., 2025), or predefined shared structures (Xiong et al., 2023). In practice, treatment assignment frequently varies across sites due to differences in norms, resources, or clinical practices. This heterogeneity implies that the global propensity score takes the form of a *weighted combination* of the site-specific scores, that is $e(X) = \sum_{k=1}^K \omega_k(X) e_k(X)$, see Appendix A.4. These weights $\{\omega_k\}_{k \in [K]}$ can be written as density ratio weights (DW):

$$e(x) = \sum_{k=1}^K \underbrace{\rho_k \frac{f_k(x)}{f(x)}}_{=\omega_k(x)} e_k(x), \quad (4)$$

where f_k and f are the covariate densities locally and globally. Similar weights are used in transportability methods that reweight data to match a target population (Han et al., 2023; 2025; Guo et al., 2024), though here the objective is to recover the global propensity score across the super-population defined by the K participating sites. Unfortunately, estimating DW ω_k in a federated setting requires modeling the f_k 's, which entails strong distributional assumptions and becomes challenging in high dimensions.

Instead, we propose to rewrite the ω_k 's as *Membership*

Weights (MW) :

$$e(x) = \sum_{k=1}^K \underbrace{\mathbb{P}(H = k \mid X = x)}_{=\omega_k(x)} e_k(x), \quad (5)$$

which represent the probability of site membership given the covariates. Unlike DW, MW do not require explicit density modeling and can be estimated directly via federated parametric (e.g., logistic regression, neural networks) or non-parametric (e.g., gradient-boosted trees) classification models, providing a flexible, communication-efficient alternative and making it the preferred choice for federated settings. We refer to Section 3.2 for more details on the federated estimation of MW and its advantages over DW.

Equations 4 and 5 enable combining locally estimated propensity scores e_k into a global propensity score using globally learned weights $\omega_k(x)$. Building on this decomposition, we define our oracle Federated IPW and AIPW estimators (Fed-(A)IPW). We define the *Oracle federated estimators* as follows:

$$\hat{\tau}_{\text{IPW}}^{\text{fed}*} = \sum_{k=1}^K \frac{n_k}{n} \hat{\tau}_{\text{IPW}}^{\text{fed}(k)}, \quad \hat{\tau}_{\text{AIPW}}^{\text{fed}*} = \sum_{k=1}^K \frac{n_k}{n} \hat{\tau}_{\text{AIPW}}^{\text{fed}(k)}, \quad (6)$$

where $\hat{\tau}_{\text{IPW}}^{\text{fed}(k)} = \frac{1}{n_k} \sum_{i=1}^{n_k} \tau_{\text{IPW}}(X_i; e)$ and $\hat{\tau}_{\text{AIPW}}^{\text{fed}(k)} = \frac{1}{n_k} \sum_{i=1}^{n_k} \tau_{\text{AIPW}}(X_i; e, \mu_1, \mu_0)$ rely on the global propensity score $e(x) = \sum_{k=1}^K \omega_k(x) e_k(x)$.

Theorem 3 (proved in Section A.5) establishes that, in the oracle setting, Fed-(A)IPW estimators attain the same efficiency as their centralized counterparts.

Theorem 3. *Under Assumptions 1, 2, and 3, the oracle federated estimators (Equation 6) are identical to the oracle centralized estimators (Equation 1).*

Theorem 4 (proved in Appendix A.6) further shows that even when local overlap (Assumption 4) holds, federated estimators have lower variance than meta-analysis estimators.

Theorem 4. *we have:*

$$\begin{aligned} \mathbb{V}[\hat{\tau}_{\text{IPW}}^*] &= \mathbb{V}[\hat{\tau}_{\text{IPW}}^{\text{fed}*}] \leq \mathbb{V}[\hat{\tau}_{\text{IPW}}^{\text{meta}*}], \\ \mathbb{V}[\hat{\tau}_{\text{AIPW}}^*] &= \mathbb{V}[\hat{\tau}_{\text{AIPW}}^{\text{fed}*}] \leq \mathbb{V}[\hat{\tau}_{\text{AIPW}}^{\text{meta}*}], \end{aligned}$$

with equality when the local propensity scores are identical across sites.

This variance reduction arises for two reasons. First, decomposing e as a weighted sum of e_k 's marginalizes over $H \mid X$, eliminating unnecessary adjustment for site membership and thereby reducing variance. Second, our federated approach improves overlap compared with meta-analysis, as formalized below.

Theorem 5. *We compare the improvement in the overlap*

$$0 \leq \mathcal{O}_{\text{global}} \leq \sum_{k=1}^K \rho_k \mathcal{O}_k.$$

Theorem 5 (proved in Appendix A.7) shows that global overlap is always at least as good as the worst local overlap. Even when local overlap holds, sites with poor overlap benefit from the federated approach because the global score $e(x)$ is more bounded away from 0 and 1 than the local scores $\{e_k(x)\}_{k \in [K]}$. Notably, sites with poor individual overlap can even improve the overall overlap of the federated population, as illustrated in the following example.

Example. Let $K = 2$ with $X_i = 1$ in both sites, $\mathbb{P}(H_i \mid X_i) = 0.5$, $e_1(X_i) = 0.99$ and $e_2(X_i) = 0.01$, leading to $e(X_i) = \sum_{k=1}^2 0.5 \times e_k(X_i) = 0.5$. Local overlaps are poor, $\mathcal{O}_1 = \mathcal{O}_2 = (0.99 \times 0.01)^{-1} \approx 101$, whereas the global overlap is $\mathcal{O}_{\text{global}} = (0.5 \times 0.5)^{-1} = 4$ —the optimal value achieved in a randomized trial with 50% treatment probability. This illustrates how heterogeneity in treatment assignments can enhance global overlap and enable more robust causal inference.

3.2. Federated Estimation

We now move beyond oracle estimators and describe how to implement our Fed-(A)IPW estimators in a practical federated learning setting. Constructing the global score propensity score requires two steps, which can be executed in parallel: each site k estimates and shares a local propensity score $\hat{e}_k(x)$; and the sites collaboratively estimate federated weights $\{\hat{\omega}_k(x)\}_{k \in [K]}$. Fed-AIPW adds a third step to train outcome models $\hat{\mu}_0, \hat{\mu}_1$ via federated learning. We detail how to estimate $\{\hat{e}_k(x), \hat{\omega}_k(x)\}_k$ and $\hat{\mu}_0, \hat{\mu}_1$ below.

Local propensity scores. Each e_k can be estimated using any probabilistic binary classifier, either parametric (e.g., logistic regression or neural networks) or non-parametric (e.g., generalized random forests, Lee et al., 2010). A key advantage of our approach is flexibility: sites can use different estimation methods tailored to local data or computational constraints. It also does *not* require Assumption 4: local scores may approach 0 or 1 provided the *global* score remains bounded away from these extremes. This, in particular, enables the integration of external control arms (FDA, 2023; EMA, 2023), where some sites have $e_k(X) = 0$ for all control patients yet still contribute to the global analysis.

Federated weights: density ratio vs. membership. Parametric density ratio weights $\hat{\omega}_k^{\text{DW}}(x) = \rho_k \frac{\hat{f}_k(x)}{\hat{f}(x)}$ can be implemented in a one-shot fashion by sharing local density parameters with the server, which then reconstructs the global mixture and computes the weights. A common

choice is to assume parametric covariate distributions (say, Gaussian), estimate $(\hat{\mu}_k, \hat{\Sigma}_k)$ locally, and transmit them once to the server. This requires to communicate $O(Kd^2)$ parameters and is highly sensitive to model misspecification—an issue that becomes critical in high dimensions. Nonparametric density estimation would relax these assumptions but is statistically inefficient and does not yet have practical federated implementations.

In contrast, $\hat{\omega}_k^{\text{MW}}(x) = \hat{\mathbb{P}}(H = k \mid X = x)$, our membership weights, can be learned with any probabilistic multi-class classifier trained federatively—for example, logistic regression for simplicity and interpretability, or neural networks to capture complex nonlinearities. Such models are readily supported by modern FL algorithms and software libraries. Using the standard FedAvg algorithm (McMahan et al., 2017) requires exchanging TKP floats (training rounds $\times$ sites $\times$ model parameters), which is feasible for models of practical size and a large number of sites. Non-parametric classifiers like random forests (Hauschild et al., 2022) and gradient-boosted trees (Li et al., 2020) have also been adapted to the federated setting.

Outcome models. To construct the doubly robust Fed-AIPW estimator, μ_0, μ_1 are trained federatively, as in Khellaf et al. (2025). As for MW, standard FL algorithms and libraries can be used to train a wide range of parametric and non-parametric supervised learning models.

Cross-fitting models in FL. Cross-fitting estimates nuisance functions on one subset of the data and evaluates them on held-out observations, thereby producing out-of-sample predictions that mitigate overfitting. For μ_w and $\{\omega_k\}_k$, implementation is straightforward: each site partitions its data into J folds, the server trains the models federatively on the corresponding training folds, and the resulting models are used to generate predictions on the held-out folds at each site. However, obtaining a cross-fitted global propensity score e additionally requires cross-fitting the local propensity scores e_k , which may become unstable when local treatment arms are small.

Statistical efficiency. We now establish an efficiency result for the empirical federated AIPW estimators, where the nuisance functions are estimated from data.

Theorem 6 (Asymptotic variance and efficiency of empirical federated estimators). *Under Assumptions 1, 2, and 4, suppose the (cross-fitted) nuisance estimators satisfy $|\hat{\mu}_w - \mu_w| = o_P(n^{-1/4})$, $|\hat{e}_k - e_k| = o_P(n^{-1/4})$ and $|\hat{\omega}_k - \omega_k| = o_P(n^{-1/4})$ for all $k \in [K]$. Because the global propensity score error is bounded by $\|\hat{e} - e\| \leq \sum_{k=1}^K \|\hat{e}_k - e_k\| + \sum_{k=1}^K \|\hat{\omega}_k - \omega_k\|$, it also achieves the $o_P(n^{-1/4})$ rate. Then, the empirical Fed-AIPW estimator attains the exact oracle variance and is locally semiparametrically efficient, i.e., $V_{\text{AIPW}}^{\text{fed}} = V_{\text{AIPW}}^*$.*

Note that if federated membership weights and local propensity models are well-specified parametric families (e.g., logistic regression), the empirical Fed-IPW estimator achieves strictly lower variance than the oracle estimator (Hirano et al., 2003).

Remark (Missing values). Our approach can naturally handle missing data. Local propensity scores can be estimated consistently with logistic regression or random forests (Jiang et al., 2020; Josse et al., 2024). For membership weights, constant imputation combined with federated random forests provides a consistent solution (Le Morvan et al., 2021), whereas density ratio weights would require adapting more complex federated EM algorithms (Dieuleveut et al., 2021; Marfoq et al., 2021). Finally, doubly robust estimators with missing data can be obtained via non-parametric federated outcome models (Mayer et al., 2020).

3.3. Accommodating Site Effects

In practice, the site membership H may affect potential outcomes through systematic differences (e.g., infrastructure, staff expertise, or measurement practices). When populations also differ across sites, ignoring such center effects can induce non-negligible confounding bias, since H becomes associated with both X and Y . In this section, we extend our framework to accommodate site effects when site-level covariates Z (e.g., center-specific resources, equipment quality etc.) are observed and adequately capture these differences. We thus propose a relaxation of Assumption 2.

Assumption 5 (Explained Site Effects).

$$\{Y(1), Y(0)\} \perp H \mid (X, Z).$$

Under Assumption 5, the effect of H on Y is fully mediated by Z , so that conditional on (X, Z) , there is no residual prognostic information carried by H . This assumption integrates naturally into our framework by incorporating Z into the membership weights decomposition of Eq. (5):

$$e(X, Z) = \sum_{k=1}^K \mathbb{P}(H = k \mid X, Z) e_k(X).$$

In practice, we learn the resulting membership weights $\hat{\omega}_k^{\text{MW}}(x, z) = \hat{\mathbb{P}}(H = k \mid X = x, Z = z)$ as described in Section 3.2, with Z included as additional covariates. Local propensity scores $e_k(X)$ remain functions of X only because Z is constant within each site. The AIPW outcome regressions are likewise extended to $\mu_w(X, Z)$.




Figure 2. Synthetic data: DGP A (top) and DGP B (bottom).

4. Experiments

4.1. Synthetic Data

We consider $K = 3$ sites and $d = 10$ covariates. Two data-generating processes (DGPs) are used. In *DGP A*, each site k independently samples $n_k = 650$ individuals from a site-specific multivariate Gaussian distribution $\mathcal{N}(\mu_k, \Sigma_k)$. In *DGP B*, a total of $n = 4000$ individuals are first drawn from a bimodal Gaussian mixture and then assigned to sites according to a multinomial logistic model based on their covariates. We vary within-site overlap to mimic different practical scenarios: *No local overlap* ($\mathcal{O}_2 = +\infty$, the second site has no treated individuals), *Poor local overlap* ($\mathcal{O}_2 \approx 10^7$), and *Good local overlap* ($\mathcal{O}_2 \approx 4.6$). The outcome models μ_1, μ_0 are shared across sites and specified as polynomial functions with interactions. For comparison, we also generate data consistent with the setting in Xiong et al. (2023) (Figure 3). All results are averaged over 800 simulation runs; full details are provided in Appendix D.

We evaluate our proposed **Fed-IPW** and **Fed-AIPW** using the **MW** weights estimated either via federated multinomial logistic regression (well specified in *DGP B*) or via a two-layer Neural Network (NN). We compare these methods against several competitors: **Fed-IPW** and **Fed-AIPW** with the alternative **DW** weights based on Gaussian density

estimation (well specified in *DGP A*); the **Centralized Oracle** (Def. 1); meta-analysis IPW/AIPW with sample-size weighting (**Meta-SW**) (Def. 3); and the one-shot inverse-variance weighted IPW estimator (**IS-IVW**) of Xiong et al. (2023), evaluated under favorable conditions with shared propensity-score parameters across sites. For all estimators, propensity scores are fit via logistic regression and outcome models via linear regression, implying misspecification in the latter.

Behavior under varying overlap. Figures 2a–2c (top: *DGP A*; bottom: *DGP B*) summarize the results under the three overlap regimes. Before discussing each setting, we highlight several general observations. First, Fed-IPW-DW is unbiased only under *DGP A*, where the Gaussian specification holds, and exhibits bias under *DGP B*. By contrast, Fed-IPW-MW adapts across data-generating processes: with logistic membership weights it is unbiased in *DGP B*, and with a more flexible neural network classifier it is unbiased for both *DGP A* and *DGP B*, reflecting the greater modeling flexibility of MW relative to DW’s restrictive density specification. Second, Fed-AIPW enjoys its doubly robust property and remains unbiased across all overlap levels by double robustness; despite misspecified linear outcome models, accurate propensity estimation ensures consistency. Finally, Fed-IPW typically at-

tains lower variance than the centralized oracle IPW, consistent with well-documented efficiency gains from using estimated propensity scores (Hirano et al., 2003).

In the *No local overlap* setting (Figure 2a), meta-analysis estimators are undefined because one site has no treated individuals. Our federated estimators remain unbiased under both DGPs, as Assumption 3 holds (global overlap $\mathcal{O}_{\text{global}} \approx 6.22$). In the *Poor local overlap* setting (Figure 2b), site 2 exhibits weak overlap, leading to bias and instability in meta-analysis estimators, including Meta-AIPW, as both the propensity scores and local outcome models are inaccurate. This issue is mitigated in the global dataset (see Figure 5 in the Appendix), allowing our federated estimators to remain reliable. In the *Good local overlap* setting (Figure 2c), all methods are unbiased, but federated estimators achieve the smallest variance.

Comparison to Xiong et al. (2023). Figure 3 considers a setting where all local propensity scores share a common subset of 5 of 10 logistic regression coefficients with data generated from *DGP B*. This setup matches the assumptions of the 1S-IVW method of Xiong et al. (2023), which relies on prior knowledge of the shared parameters to aggregate them—an assumption not required by our method. Fed-MW remains unbiased and attains the lowest variance, matching that of 1S-IVW, even without access to the additional information about shared parameters.

Handling site effects. We assess the benefit of adjusting for site effects when site-level covariates Z are observed (Section 3.3). We simulate $K = 30$ sites, with heterogeneity arising both from diverse distributions of individual-level covariates X (via covariate shift induced by $H \sim \text{Logit}(\theta^\top X)$) and from site-level variation in outcomes captured by Z . Outcomes follow $Y = X^\top \beta + Z^\top \delta + \tau W + \varepsilon$, so that H affects Y only through Z . Federated methods incorporate Z only in the membership weights, while meta-analysis adjusts only for X as it naturally handles site effects. As shown in Table 1, incorporating Z into our method eliminates confounding due to both covariate shift ($H \rightarrow X$) and mediated site effects ($H \rightarrow Z \rightarrow Y$). As a result, our Federated IPW/AIPW markedly improves MSE, especially for small n_k , where meta-analysis suffers from unstable local nuisance estimation.

Additional simulations provided in Appendix F examine scenarios with more sites (Table 5), nonparametric estimation (Table 6), and local propensity model misspecification (Table 7), further confirming the robustness of our approach.

4.2. Real Data

We analyze the multi-site Traumabase cohort (Mayer et al., 2020; Colnet et al., 2024) to estimate the effect of tranex-

Estimator	$n_k=20$	$n_k=60$	$n_k=100$
IPW Oracle ($e^*(X, Z)$)	0.853	0.443	0.091
IPW Meta-SW ($\hat{e}(X)$)	5.176	0.777	0.590
IPW Fed-MW (Logit, $\hat{e}(X, Z)$)	0.427	0.132	0.043
IPW Fed-MW (Logit, $\hat{e}(X)$)	2.001	1.384	0.327
IPW Fed-MW (NN, $\hat{e}(X, Z)$)	1.083	0.465	0.206
IPW Fed-MW (NN, $\hat{e}(X)$)	2.100	1.418	0.356
AIPW Oracle ($e^*(X, Z)$)	0.015	0.005	0.002
IPW Meta-SW ($\hat{e}(X)$)	10.673	0.346	0.084
AIPW Fed-MW (Logit, $\hat{e}(X, Z)$)	0.022	0.004	0.002
AIPW Fed-MW (Logit, $\hat{e}(X)$)	2.026	1.348	0.323
AIPW Fed-MW (NN, $\hat{e}(X, Z)$)	0.020	0.004	0.002
AIPW Fed-MW (NN, $\hat{e}(X)$)	2.043	1.349	0.315

Table 1. Mean Squared Error (MSE) of federated and meta-analysis estimators under Assumption 5.

amic acid on mortality across $K = 14$ centers. The local datasets are highly imbalanced in sizes (106 to 2,092 patients) and treatment arms (e.g., site 11: 4 treated vs. 121 controls); there are 17 covariates and $n = 8,248$ patients in total, of whom 638 were treated. Covariates are standardized federatively by sharing site means and variances. We focus on AIPW estimators: local propensities e_k are estimated using the R package `grf`’s `probability_forest` function (Tibshirani et al., 2018) on large sites, and logistic regressions on smaller ones; outcome models μ_1, μ_0 are trained with FedAvg logistic regression (5,000 rounds, 1 local epoch, step size $\eta = 0.1$); MW are learned with a FedAvg neural network (NN, one hidden layer, 128 units) and a federated logistic regression. We also report a centralized random forest (RF) MW as a strong benchmark. DW are based on Gaussian density estimation. Competitors include Meta-SW (same local models, computed only on sites with enough treated units—where the number of treated observations exceeds the covariate dimension) and a centralized AIPW benchmark obtained using `probability_forest` function on the pooled data. Empirical confidence intervals are constructed from 150 bootstrap resamples.

Figure 4 shows that our federated MW estimators closely match the centralized benchmark and exhibit lower variance than Meta-SW, which in this application departs from the centralized estimate. Among federated methods, Fed-MW are closest to the nonparametric centralized estimators—reflecting their more flexible modeling of propensities and membership—whereas Fed-DW is less stable, likely due to noisy per-site covariance estimates for a 17-dimensional Gaussian ($\approx 17^2$ parameters) in small sites.




Figure 3. Comparison to Xiong et al. (2023) (IPW estimators).




Figure 4. AIPW estimates on Traumabase.

5. Conclusion, Limitations and Future work

We propose a theoretically grounded framework for federated causal inference that leverages membership weights to construct valid causal estimates across silos. These weights, estimated via flexible parametric or nonparametric models, improve overlap and yield more stable ATE estimates. Our framework accommodates arbitrary heterogeneity in local propensity score functions, does not require local overlap to hold, enabling the inclusion of small sites and external control arms, without sharing individual-level data. Our approach can also be naturally extended to handle some form of site effects. Although sufficiently large per-site datasets are still needed—particularly in high-dimensional settings—to learn local propensity scores accurately, our approach is especially well suited to a moderate number of large silos, where federated learning most effectively increases effective sample size.

A promising avenue for future work is to extend our framework to CATE estimation. While the ATE remains central in econometrics, biomedicine, and public policy, considering CATE is important for moving towards personalization. Our work lays a foundation for federated CATE estimation: most learners (T-, S-, X-, and R-learners) rely on nuisance quantities such as propensity scores, and our federated estimation framework could be incorporated into these methods, although further work is needed to assess its formal properties and practical performance.

Impact Statement

This paper presents work whose goal is to advance the field of machine learning. There are many potential societal consequences of our work, none of which we feel must be specifically highlighted here.

References

Heejung Bang and James M Robins. Doubly robust estimation in missing data and causal inference models. *Biometrics*, 61(4):962–973, 2005.

Clément Berenfeld, Ahmed Boughdiri, Bénédicte Colnet, Wouter AC van Amsterdam, Aurélien Bellet, Rémi Khellaf, Erwan Scornet, and Julie Josse. Causal meta-analysis: Rethinking the foundations of evidence-based medicine. *arXiv preprint arXiv:2505.20168*, 2025.

Danielle L Burke, Joie Ensor, and Richard D Riley. Meta-analysis using individual participant data: one-stage and two-stage approaches, and why they may differ. *Statistics in medicine*, 36(5):855–875, 2017.

Victor Chernozhukov, Denis Chetverikov, Mert Demirer, Esther Duflo, Christian Hansen, Whitney Newey, and James Robins. Double/debiased machine learning for treatment and structural parameters, 2018.

Oscar Clivio, David Bruns-Smith, Avi Feller, and Christopher C Holmes. Towards principled representation learning to improve overlap in treatment effect estimation. In *9th Causal Inference Workshop at UAI 2024*, 2024.

Bénédicte Colnet, Imke Mayer, Guanhua Chen, Awa Dieng, Ruohong Li, Gaël Varoquaux, Jean-Philippe Vert, Julie Josse, and Shu Yang. Causal inference methods for combining randomized trials and observational studies: a review. *Statistical science*, 39(1):165–191, 2024.

A Philip Dawid. Conditional independence in statistical theory. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 41(1):1–15, 1979.

Aymeric Dieuleveut, Gersende Fort, Eric Moulines, and Geneviève Robin. Federated expectation maximization with heterogeneity mitigation and variance reduction. *CoRR*, abs/2111.02083, 2021. URL <https://arxiv.org/abs/2111.02083>.

EMA. Reflection paper on establishing efficacy based on single arm trials submitted as pivotal evidence in a marketing authorisation, 2023. European Medicines Agency.

EMA. ICH E9 Statistical Principles for Clinical Trials: Scientific Guideline, 2024.

- FDA. Adjusting for covariates in randomized clinical trials for drugs and biological products, 2023.
- FDA. Considerations for the design and conduct of externally controlled trials for drug and biological products. <https://www.fda.gov/media/164960/download>, 2023. Center for Drug Evaluation and Research and Center for Biologics Evaluation and Research and Oncology Center of Excellence, Food and Drug Administration.
- David A Grimes and Kenneth F Schulz. Bias and causal associations in observational research. *The lancet*, 359 (9302):248–252, 2002.
- Tianyu Guo, Sai Praneeth Karimireddy, and Michael I Jordan. Collaborative heterogeneous causal inference beyond meta-analysis. *arXiv preprint arXiv:2404.15746*, 2024.
- Zijian Guo, Xiudi Li, Larry Han, and Tianxi Cai. Robust inference for federated meta-learning. *Journal of the American Statistical Association*, pp. 1–16, 2025.
- Larry Han, Zhu Shen, and Jose Zubizarreta. Multiply robust federated estimation of targeted average treatment effects. *Advances in Neural Information Processing Systems*, 36:70453–70482, 2023.
- Larry Han, Jue Hou, Kelly Cho, Rui Duan, and Tianxi Cai. Federated adaptive causal estimation (face) of target treatment effects. *Journal of the American Statistical Association*, (just-accepted):1–25, 2025.
- Anne-Christin Hauschild, Marta Lemanczyk, Julian Matschinske, Tobias Frisch, Olga Zolotareva, Andreas Holzinger, Jan Baumbach, and Dominik Heider. Federated random forests can improve local performance of predictive models for various healthcare applications. *Bioinformatics*, 38(8):2278–2286, 2022.
- Miguel A Hernán. The c-word: scientific euphemisms do not improve causal inference from observational data. *American journal of public health*, 108(5):616–619, 2018.
- Miguel A Hernán and James M Robins. Estimating causal effects from epidemiological data. *Journal of Epidemiology & Community Health*, 60(7):578–586, 2006.
- Keisuke Hirano, Guido W Imbens, and Geert Ridder. Efficient estimation of average treatment effects using the estimated propensity score. *Econometrica*, 71(4):1161–1189, 2003.
- Wei Jiang, Julie Josse, Marc Lavielle, TraumaBase Group, et al. Logistic regression with missing covariates—parameter estimation, model selection and prediction within a joint-modeling framework. *Computational Statistics & Data Analysis*, 145:106907, 2020.
- Julie Josse, Jacob M Chen, Nicolas Prost, Gaël Varoquaux, and Erwan Scornet. On the consistency of supervised learning with missing values. *Statistical Papers*, 65(9): 5447–5479, 2024.
- Peter Kairouz, H Brendan McMahan, Brendan Avent, Aurélien Bellet, Mehdi Bennis, Arjun Nitin Bhagoji, Kallista Bonawitz, Zachary Charles, Graham Cormode, Rachel Cummings, et al. Advances and open problems in federated learning. *Foundations and trends® in machine learning*, 14(1–2):1–210, 2021.
- Ahmed Khaled, Konstantin Mishchenko, and Peter Richtárik. Tighter theory for local SGD on identical and heterogeneous data. In *AISTATS*, 2020.
- Rémi Khellaf, Aurélien Bellet, and Julie Josse. Federated Causal Inference: Multi-Study ATE Estimation beyond Meta-Analysis. In *AISTATS*, 2025.
- Marine Le Morvan, Julie Josse, Erwan Scornet, and Gaël Varoquaux. What’s a good imputation to predict with missing values? *Advances in Neural Information Processing Systems*, 34:11530–11540, 2021.
- Brian K Lee, Justin Lessler, and Elizabeth A Stuart. Improving propensity score weighting using machine learning. *Statistics in medicine*, 29(3):337–346, 2010.
- Lihua Lei and Peng Ding. Regression adjustment in completely randomized experiments with a diverging number of covariates. *Biometrika*, 108(4):815–828, 2021.
- Fan Li, Kari Lock Morgan, and Alan M. Zaslavsky and. Balancing covariates via propensity score weighting. *Journal of the American Statistical Association*, 113 (521):390–400, 2018a. doi: 10.1080/01621459.2016.1260466. URL <https://doi.org/10.1080/01621459.2016.1260466>.
- Fan Li, Laine E Thomas, and Fan Li. Addressing extreme propensity scores via the overlap weights. *American Journal of Epidemiology*, 188(1):250–257, 09 2018b. ISSN 0002-9262. doi: 10.1093/aje/kwy201. URL <https://doi.org/10.1093/aje/kwy201>.
- Qinbin Li, Zeyi Wen, and Bingsheng He. Practical federated gradient boosting decision trees. In *AAAI*, pp. 4642–4649, 2020.
- Xiang Li, Kaixuan Huang, Wenhao Yang, Shusen Wang, and Zhihua Zhang. On the convergence of fedavg on non-iid data. *arXiv preprint arXiv:1907.02189*, 2019.

- Othmane Marfoq, Giovanni Neglia, Aurélien Bellet, Laetitia Kameni, and Richard Vidal. Federated Multi-Task Learning under a Mixture of Distributions. In *NeurIPS*, 2021.
- Imke Mayer, Erik Sverdrup, Tobias Gauss, Jean-Denis Moyer, Stefan Wager, and Julie Josse. Doubly robust treatment effect estimation with missing attributes. *The Annals of Applied Statistics*, 14(3):1409–1431, 2020.
- Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Agüera y Arcas. Communication-efficient learning of deep networks from decentralized data. In *Artificial intelligence and statistics*, pp. 1273–1282. PMLR, 2017.
- Fernando P. Polack, Stephen J. Thomas, Nicholas Kitchin, Judith Absalon, Alejandra Gurtman, Stephen Lockhart, John L. Perez, Gonzalo Pérez Marc, Edson D. Moreira, Cristiano Zerbini, Ruth Bailey, Kena A. Swanson, Satrajit Roychoudhury, Kenneth Koury, Ping Li, Warren V. Kalina, David Cooper, Robert W. Frenc, Laura L. Hammitt, Özlem Türeci, Haylene Nell, Axel Schaefer, Serhat Ünal, Dina B. Tresnan, Susan Mather, Philip R. Dormitzer, Uğur Şahin, Kathrin U. Jansen, and William C. Gruber. Safety and efficacy of the bnt162b2 mrna covid-19 vaccine. *New England Journal of Medicine*, 383(27):2603–2615, 2020. doi: 10.1056/NEJMoa2034577. URL <https://www.nejm.org/doi/full/10.1056/NEJMoa2034577>.
- Richard D Riley, Joie Ensor, Miriam Hattle, Katerina Papadimitropoulou, and Tim P Morris. Two-stage or not two-stage? that is the question for ipd meta-analysis projects. *Research Synthesis Methods*, 14(6):903–910, 2023.
- Sarah E Robertson, Jon A Steingrimsdóttir, Nina R Joyce, Elizabeth A Stuart, and Issa J Dahabreh. Center-specific causal inference with multicenter trials: reinterpreting trial evidence in the context of each participating center. *arXiv preprint arXiv:2104.05905*, 2021.
- James Robins. A new approach to causal inference in mortality studies with a sustained exposure period—application to control of the healthy worker survivor effect. *Mathematical modelling*, 7(9-12):1393–1512, 1986.
- Paul R Rosenbaum and Donald B Rubin. The central role of the propensity score in observational studies for causal effects. *Biometrika*, 70(1):41–55, 1983.
- Donald B Rubin. Estimating causal effects of treatments in randomized and nonrandomized studies. *Journal of educational psychology*, 66(5):688–701, 1974. ISSN 0022-0663.
- Jerzy Splawa-Neyman. On the Application of Probability Theory to Agricultural Experiments. Essay on Principles. Section 9. *Statistical Science*, 5(4):465 – 472, 1990. doi: 10.1214/ss/1177012031. URL <https://doi.org/10.1214/ss/1177012031>.
- Sebastian U. Stich. Local sgd converges fast and communicates little. In *ICLR*, 2019.
- Julie Tibshirani, Susan Athey, Rina Friedberg, Vitor Hadad, David Hirshberg, Luke Miner, Erik Sverdrup, Stefan Wager, Marvin Wright, and Maintainer Julie Tibshirani. Package ‘grf’. *Comprehensive R Archive Network*, 2018.
- Tyler J VanderWeele. Principles of confounder selection. *European journal of epidemiology*, 34:211–219, 2019.
- Thanh Vinh Vo, Young Lee, Trong Nghia Hoang, and Tze-Yun Leong. Bayesian federated estimation of causal effects from observational data. In James Cussens and Kun Zhang (eds.), *Proceedings of the Thirty-Eighth Conference on Uncertainty in Artificial Intelligence*, volume 180 of *Proceedings of Machine Learning Research*, pp. 2024–2034. PMLR, 01–05 Aug 2022. URL <https://proceedings.mlr.press/v180/vo22a.html>.
- Stefan Wager. Causal inference: A statistical learning approach, 2024.
- Ruoxuan Xiong, Allison Koenecke, Michael Powell, Zhu Shen, Joshua T Vogelstein, and Susan Athey. Federated causal inference in heterogeneous observational data. *Statistics in Medicine*, 42(24):4418–4439, 2023.
- Changchang Yin, Hong-You Chen, Wei-Lun Chao, and Ping Zhang. Federated inverse probability treatment weighting for individual treatment effect estimation. *arXiv preprint arXiv:2503.04946*, 2025.

A. Proofs

A.1. ATE Identification and Sufficient Conditions

Proof. Under Assumptions 1 and 2, the variables in X form a sufficient adjustment set. We identify the ATE as follows:

$$\begin{aligned}
 \tau &= \mathbb{E}[Y(1) - Y(0)] \\
 &= \mathbb{E}[\mathbb{E}[Y(1) | X, H] - \mathbb{E}[Y(0) | X, H]] && \text{Law of Iterated Expectations} \\
 &= \mathbb{E}[\mathbb{E}[Y(1) | X, H, W = 1] - \mathbb{E}[Y(0) | X, H, W = 0]] && \text{Assumption 1} \\
 &= \mathbb{E}[\mathbb{E}[Y(1) | X, W = 1] - \mathbb{E}[Y(0) | X, W = 0]] && \text{Assumption 2 and Lemma 4.3 (Dawid, 1979)} \\
 &= \mathbb{E}[\mathbb{E}[Y | X, W = 1] - \mathbb{E}[Y | X, W = 0]] && \text{Consistency} \\
 &= \mathbb{E}\left[\frac{WY}{e(X)} - \frac{(1-W)Y}{1-e(X)}\right] && \text{Assumption 3}
 \end{aligned}$$

□

A.2. Proof of Theorem 1

Proof. **Unbiasedness.** We show the unbiasedness of the first term of the IPW estimator. The derivation for the second term is symmetric.

$$\begin{aligned}
 \mathbb{E}\left[\frac{WY}{e(X)}\right] &= \mathbb{E}\left[\mathbb{E}\left[\frac{WY(1)}{e(X)} \middle| X, H\right]\right] && \text{SUTVA} \\
 &= \mathbb{E}\left[\frac{\mathbb{E}[W | X, H] \mathbb{E}[Y(1) | X, H]}{e(X)}\right] && \text{Assumption 1} \\
 &= \mathbb{E}\left[\frac{e_H(X) \mathbb{E}[Y(1) | X, H]}{e(X)}\right] && \text{Definition of local score } e_H(X) \\
 &= \mathbb{E}\left[\frac{e_H(X) \mathbb{E}[Y(1) | X]}{e(X)}\right] && \text{Assumption 2} \\
 &= \mathbb{E}_X\left[\frac{\mathbb{E}[Y(1) | X]}{e(X)} \mathbb{E}_H[e_H(X) | X]\right] && \text{Iterated Expectation over } H \\
 &= \mathbb{E}_X\left[\frac{\mathbb{E}[Y(1) | X]}{e(X)} \sum_{k=1}^K \mathbb{P}(H = k | X) e_k(X)\right] \\
 &= \mathbb{E}_X\left[\frac{\mathbb{E}[Y(1) | X]}{e(X)} e(X)\right] && \text{Definition of global } e(X) \\
 &= \mathbb{E}[\mathbb{E}[Y(1) | X]] = \mathbb{E}[Y(1)]
 \end{aligned}$$

Note that the step applying Assumption 2 relies on the consequence that $\mathbb{E}[Y(1) | X, H] = \mathbb{E}[Y(1) | X]$ (which is derived from Assumption 2 combined with Assumption 1).

Similarly, we have $\mathbb{E}\left[\frac{(1-W)Y}{1-e(X)}\right] = \mathbb{E}[Y(0)]$, so that $\mathbb{E}[\hat{\tau}^{\text{IPW}^*}] = \tau$, which concludes the proof of unbiasedness of the oracle multi-site centralized IPW estimator.

For the AIPW estimator, we note that $\mathbb{E}[W(Y - \mu_1(X))/e(X)] = 0$ following the same logic as above (replacing $Y(1)$ with the residual $Y(1) - \mu_1(X)$ which has mean zero given X). Thus:

$$\begin{aligned}
 \mathbb{E}[\hat{\tau}^{\text{AIPW}^*}] &= \mathbb{E}\left[\mu_1(X) - \mu_0(X) + \frac{W(Y - \mu_1(X))}{e(X)} - \frac{(1-W)(Y - \mu_0(X))}{1-e(X)}\right] \\
 &= \mathbb{E}[\mu_1(X) - \mu_0(X)] + 0 - 0 \\
 &= \mathbb{E}[\tau(X)] = \tau
 \end{aligned}$$

Variance. As we consider uniformly bounded potential outcomes, that is $\forall w \in \{0, 1\}, \exists (L, U) \in \mathbb{R}^2, L < Y(w) < U$, and Assumption 3, we have that $\mathbb{E} \left[\frac{Y_i(1)^2}{e(X_i)} \right] < \infty$ and $\mathbb{E} \left[\frac{Y_i(0)^2}{1-e(X_i)} \right] < \infty$, so the quantities that follow are well defined:

$$\begin{aligned}
 \mathbb{V} \left[\hat{\tau}^{\text{IPW}^*} \right] &= \mathbb{V} \left[\frac{1}{n} \sum_{i=1}^n \left(\frac{W_i Y_i}{e(X_i)} - \frac{(1-W_i)Y_i}{1-e(X_i)} \right) \right] \\
 &= \frac{1}{n^2} \sum_{i=1}^n \mathbb{V} \left[\frac{W_i Y_i}{e(X_i)} - \frac{(1-W_i)Y_i}{1-e(X_i)} \right] && \text{i.i.d. observations} \\
 &= \frac{1}{n^2} \sum_{i=1}^n \left(\mathbb{E} \left[\left(\frac{W_i Y_i}{e(X_i)} - \frac{(1-W_i)Y_i}{1-e(X_i)} \right)^2 \right] - \tau^2 \right) && \text{unbiasedness} \\
 &= \frac{1}{n^2} \sum_{i=1}^n \left(\mathbb{E} \left[\left(\frac{W_i Y_i}{e(X_i)} \right)^2 \right] + \mathbb{E} \left[\left(\frac{(1-W_i)Y_i}{1-e(X_i)} \right)^2 \right] - 2\mathbb{E} [0] - \tau^2 \right) \\
 &= \frac{1}{n^2} \sum_{i=1}^n \left(\mathbb{E} \left[\left(\frac{W_i Y_i}{e(X_i)} \right)^2 \right] + \mathbb{E} \left[\left(\frac{(1-W_i)Y_i}{1-e(X_i)} \right)^2 \right] - \tau^2 \right) \\
 \\
 \mathbb{E} \left[\left(\frac{W_i Y_i}{e(X_i)} \right)^2 \right] &= \mathbb{E} \left[\mathbb{E} \left[\left(\frac{W_i Y_i}{e(X_i)} \mid X_i \right)^2 \right] \right] \\
 &= \mathbb{E} \left[\frac{\mathbb{E} \left[(W_i Y_i \mid X_i)^2 \right]}{e(X_i)^2} \right] \\
 &= \mathbb{E} \left[\frac{\mathbb{E} [W_i \mid X_i] \mathbb{E} [Y_i^2 \mid X_i, W_i = 1]}{e(X_i)^2} \right] && \text{Assumption 1} \\
 &= \mathbb{E} \left[\frac{\mathbb{E} [Y_i(1)^2 \mid X_i]}{e(X_i)} \right] && \text{SUTVA} \\
 &= \mathbb{E} \left[\frac{Y_i(1)^2}{e(X_i)} \right]
 \end{aligned}$$

Similarly, $\mathbb{E} \left[\left(\frac{(1-W_i)Y_i}{1-e(X_i)} \right)^2 \right] = \mathbb{E} \left[\frac{Y_i(0)^2}{1-e(X_i)} \right]$.

Then, the variance of the oracle multi-site IPW estimator is

$$\mathbb{V} \left[\hat{\tau}^{\text{IPW}^*} \right] = \frac{1}{n} \left(\mathbb{E} \left[\frac{Y_i(1)^2}{e(X_i)} \right] + \mathbb{E} \left[\frac{Y_i(0)^2}{1-e(X_i)} \right] - \tau^2 \right).$$

For the variance of the oracle multi-site centralized AIPW, we first notice that since $\mathbb{E} [Y_i(w) - \mu_w(X_i) \mid X_i] = 0$ for treatment w , we have

$$\begin{aligned}
 A &= \text{Cov} \left(\tau(X_i), \frac{W_i(Y_i - \mu_1(X_i))}{e(X_i)} - \frac{(1-W_i)(Y_i - \mu_0(X_i))}{1-e(X_i)} \right) \\
 &= \mathbb{E} \left[\tau(X_i) \left(\frac{W_i(Y_i - \mu_1(X_i))}{e(X_i)} - \frac{(1-W_i)(Y_i - \mu_0(X_i))}{1-e(X_i)} \right) \right] \\
 &= \mathbb{E} \left[\tau(X_i) \mathbb{E} \left[\frac{W_i(Y_i - \mu_1(X_i))}{e(X_i)} \mid X_i \right] \right] - \mathbb{E} \left[\tau(X_i) \mathbb{E} \left[\frac{(1-W_i)(Y_i - \mu_0(X_i))}{1-e(X_i)} \mid X_i \right] \right] \\
 &= \mathbb{E} \left[\tau(X_i) \frac{e(X_i) \mathbb{E} [Y_i(1) - \mu_1(X_i) \mid X_i]}{e(X_i)} \right] - \mathbb{E} \left[\tau(X_i) \frac{(1-e(X_i)) \mathbb{E} [Y_i(0) - \mu_0(X_i) \mid X_i]}{1-e(X_i)} \right] \\
 &= 0
 \end{aligned}$$

Then,

$$\begin{aligned}
 \mathbb{V} [\hat{\tau}^{\text{AIPW}^*}] &= \mathbb{V} \left[\frac{1}{n} \sum_{i=1}^n \left(\mu_1(X_i) - \mu_0(X_i) + \frac{W_i(Y_i - \mu_1(X_i))}{e(X_i)} - \frac{(1 - W_i)(Y_i - \mu_0(X_i))}{1 - e(X_i)} \right) \right] \\
 &= \frac{1}{n^2} \sum_{i=1}^n \left(\mathbb{V} [\tau(X_i)] + \mathbb{V} \left[\frac{W_i(Y_i - \mu_1(X_i))}{e(X_i)} \right] + \mathbb{V} \left[\frac{(1 - W_i)(Y_i - \mu_0(X_i))}{1 - e(X_i)} \right] + 2A \right) \\
 &= \frac{1}{n^2} \sum_{i=1}^n \left(\mathbb{V} [\tau(X_i)] + \mathbb{E} \left[\left(\frac{Y - \mu_1(X_i)}{e(X_i)} \right)^2 \right] + \mathbb{E} \left[\left(\frac{Y - \mu_0(X_i)}{1 - e(X_i)} \right)^2 \right] \right) \\
 &= \frac{1}{n} \left(\mathbb{V} [\tau(X_i)] + \mathbb{E} \left[\left(\frac{Y - \mu_1(X_i)}{e(X_i)} \right)^2 \right] + \mathbb{E} \left[\left(\frac{Y - \mu_0(X_i)}{1 - e(X_i)} \right)^2 \right] \right)
 \end{aligned}$$

Because these variances are well defined from what precedes, the Central Limit Theorem can be applied to $\hat{\tau}^{\text{IPW}^*}$ and $\hat{\tau}^{\text{AIPW}^*}$, which gives the result in Theorem 1. $\square$

A.3. Proof of Theorem 2

Proof. **Unbiasedness.** We first prove the unbiasedness of the local IPW estimators:

$$\begin{aligned}
 \mathbb{E} [\hat{\tau}_k^{\text{IPW}^*} \mid H_i = k] &= \mathbb{E} \left[\frac{W_i Y_i}{e_k(X_i)} - \frac{(1 - W_i) Y_i}{1 - e_k(X_i)} \mid H_i = k \right] && \text{i.i.d.} \\
 &= \mathbb{E} \left[\frac{W_i Y_i(1)}{e_k(X_i)} - \frac{(1 - W_i) Y_i(0)}{1 - e_k(X_i)} \mid H_i = k \right] && \text{SUTVA} \\
 &= \mathbb{E} \left[\mathbb{E} \left[\frac{W_i Y_i(1)}{e_k(X_i)} \mid H_i = k, X_i \right] - \mathbb{E} \left[\frac{(1 - W_i) Y_i(0)}{1 - e_k(X_i)} \mid X_i, H_i = k, X_i \right] \right] \\
 &= \mathbb{E} \left[\frac{\mathbb{E} [W_i \mid H_i = k, X_i] \mathbb{E} [Y_i(1) \mid H_i = k, X_i]}{e_k(X_i)} \right. \\
 &\quad \left. - \frac{\mathbb{E} [(1 - W_i) \mid H_i = k, X_i] \mathbb{E} [Y_i(0) \mid H_i = k, X_i]}{1 - e_k(X_i)} \right] && \text{Ass. 1} \\
 &= \mathbb{E} [\mathbb{E} [Y_i(1) - Y_i(0) \mid H_i = k, X_i]] \\
 &= \mathbb{E} [Y_i(1) - Y_i(0) \mid H_i = k] \\
 &= \tau_k
 \end{aligned}$$

This yields

$$\begin{aligned}
 \mathbb{E} [\hat{\tau}^{\text{meta-IPW}^*}] &= \mathbb{E} \left[\sum_{k=1}^K \frac{n_k}{n} \hat{\tau}_k^{\text{IPW}^*} \right] \\
 &= \sum_{k=1}^K \mathbb{E} \left[\frac{n_k}{n} \mathbb{E} [\hat{\tau}_k^{\text{IPW}^*} \mid H_i = k] \right] \\
 &= \sum_{k=1}^K \mathbb{E} \left[\frac{n_k}{n} \right] \tau_k \\
 &= \sum_{k=1}^K \rho_k \tau_k \\
 &= \tau
 \end{aligned}$$

Variance. For the variance of the local ATEs, we follow the same steps as in the previous proof which yields

$$\begin{aligned}\mathbb{V}\left[\hat{\tau}_k^{\text{IPW}^*} \mid H = k\right] &= \frac{1}{n_k} \left(\mathbb{E}\left[\frac{Y_i(1)^2}{e_k(X_i)} \mid H_i = k\right] - \mathbb{E}\left[\frac{Y_i(0)^2}{1 - e_k(X_i)} \mid H_i = k\right] - \tau_k^2 \right) \\ &= \frac{1}{n_k} V_{k,i}\end{aligned}$$

with $V_{k,i} = \mathbb{E}\left[\frac{Y_i(1)^2}{e_k(X_i)} \mid H_i = k\right] - \mathbb{E}\left[\frac{Y_i(0)^2}{1 - e_k(X_i)} \mid H_i = k\right] - \tau_k^2$. Finally, by Lemma 2, we have

$$\begin{aligned}\mathbb{V}\left[\hat{\tau}^{\text{meta-IPW}^*}\right] &= \mathbb{E}\left[\mathbb{V}\left[\hat{\tau}^{\text{meta-IPW}^*} \mid H = k\right]\right] + \mathbb{V}\left[\mathbb{E}\left[\hat{\tau}^{\text{meta-IPW}^*} \mid H = k\right]\right] \\ &= \frac{1}{n} \sum_{k=1}^K \rho_k V_{k,i} + \frac{1}{n} \mathbb{V}[\tau_H]\end{aligned}$$

The proof for the meta AIPW estimator follows the same steps. □

A.4. Decomposition of the global propensity score

By the law of total probabilities,

$$\begin{aligned}e(x) &= \mathbb{P}(W_i = 1 \mid X_i) \\ &= \sum_{k=1}^K \mathbb{P}(W_i = 1 \cap H_i = k \mid X_i = x) \\ &= \sum_{k=1}^K \mathbb{P}(H_i = k \mid X_i) \mathbb{P}(W_i = 1 \mid X_i = x, H_i = k) \\ &= \sum_{k=1}^K \mathbb{P}(H_i = k) \frac{P(X_i \mid H_i = k)}{P(X_i)} e_k(X_i) & (\text{DW}) \\ &= \sum_{k=1}^K \mathbb{P}(H_i = k \mid X_i) e_k(X_i) & (\text{MW})\end{aligned}$$

A.5. Proof of Theorem 3

Proof.

$$\begin{aligned}\hat{\tau}_{\text{IPW}}^{\text{fed}^*} &= \sum_{k=1}^K \frac{n_k}{n} \hat{\tau}_{\text{IPW}}^{\text{fed}(k)} \\ &= \sum_{k=1}^K \frac{n_k}{n} \left(\frac{1}{n_k} \sum_{i=1}^{n_k} \left(\frac{W_i Y_i}{e(X_i)} - \frac{(1 - W_i) Y_i}{1 - e(X_i)} \right) \right) \\ &= \frac{1}{n} \sum_{i=1}^n \left(\frac{W_i Y_i}{e(X_i)} - \frac{(1 - W_i) Y_i}{1 - e(X_i)} \right) \\ &= \hat{\tau}^{\text{IPW}^*}\end{aligned}$$

We prove similarly that $\hat{\tau}_{\text{AIPW}}^{\text{fed}^*} = \hat{\tau}^{\text{AIPW}^*}$. □

A.6. Proof of Theorem 4

We start by two technical lemmas.

Lemma 1. Under Assumptions 1 and 2 we have,

$$\mathbb{V}[\tau(X_i)] = \sum_{k=1}^K \rho_k \mathbb{V}[\tau(X_i) | H_i = k] + \mathbb{V}[\tau_{H_i}]$$

Proof.

$$\begin{aligned} \mathbb{V}[\tau(X_i)] &= \mathbb{E}[\mathbb{V}[\tau(X_i) | H_i = k]] + \mathbb{V}[\mathbb{E}[\tau(X_i) | H_i = k]] \\ &= \sum_{k=1}^K \rho_k \mathbb{V}[\tau(X_i) | H_i = k] + \mathbb{V}\left[\sum_{k=1}^K \mathbb{1}_{[H_i=k]} \tau_{H_i}\right] \\ &= \sum_{k=1}^K \rho_k \mathbb{V}[\tau(X_i) | H_i = k] + \mathbb{V}\left[\tau_{H_i} \sum_{k=1}^K \mathbb{1}_{[H_i=k]}\right] \\ &= \sum_{k=1}^K \rho_k \mathbb{V}[\tau(X_i) | H_i = k] + \mathbb{V}[\tau_{H_i}] \end{aligned}$$

□

Lemma 2. For the general form of meta-analysis estimator $\hat{\tau}^{\text{meta}} = \sum_{k=1}^K \frac{n_k}{n} \hat{\tau}_k$, we have

$$\mathbb{V}[\hat{\tau}^{\text{meta}}] = \frac{1}{n} \sum_{k=1}^K \rho_k V_k + \frac{1}{n} \mathbb{V}[\tau_{H_i}]$$

where $V_k = \mathbb{V}[\hat{\tau}_k | H_i = k]$ is the within-site variance of the ATE estimator in site k and $\mathbb{V}[\tau_{H_i}] = \mathbb{V}[\mathbb{E}[Y_i(1) - Y_i(0) | H_i]]$ is the between-sites variance of the local ATEs.

Proof.

$$\begin{aligned} \mathbb{V}[\hat{\tau}^{\text{meta}}] &= \mathbb{V}\left[\sum_{k=1}^K \frac{n_k}{n} \hat{\tau}_k\right] \\ &= \mathbb{E}\left[\mathbb{V}\left[\sum_{k=1}^K \frac{n_k}{n} \hat{\tau}_k \mid H_1, \dots, H_n\right]\right] + \mathbb{V}\left[\mathbb{E}\left[\sum_{k=1}^K \frac{n_k}{n} \hat{\tau}_k \mid H_1, \dots, H_n\right]\right] \\ &= \mathbb{E}\left[\sum_{k=1}^K \frac{n_k}{n^2} V_k\right] + \mathbb{V}\left[\sum_{k=1}^K \frac{n_k}{n} \tau_k\right] = \frac{1}{n} \sum_{k=1}^K \rho_k V_k + \mathbb{V}\left[\frac{1}{n} \sum_{k=1}^K \sum_{i=1}^n \mathbb{1}_{[H_i=k]} \tau_{H_i}\right] \\ &= \frac{1}{n} \sum_{k=1}^K \rho_k V_k + \mathbb{V}\left[\frac{1}{n} \sum_{i=1}^n \tau_{H_i} \sum_{k=1}^K \mathbb{1}_{[H_i=k]}\right] = \frac{1}{n} \sum_{k=1}^K \rho_k V_k + \mathbb{V}\left[\frac{1}{n} \sum_{i=1}^n \tau_{H_i}\right] \\ &= \frac{1}{n} \sum_{k=1}^K \rho_k V_k + \frac{1}{n} \mathbb{V}[\tau_H], \end{aligned}$$

□

We can now prove Theorem 4.

Proof. We begin with IPW estimators, and then move to AIPW.

IPW. First, applying Jensen's inequality with the strictly convex function $t \mapsto \frac{1}{t}$ in $]0; 1[$ and summing-to-one weights $\omega_k(X) = \mathbb{P}(H_i = k \mid X_i = X)$, we have

$$\mathbb{E} \left[\frac{Y_i(1)^2}{e(X_i)} \right] < \mathbb{E} \left[\sum_{k=1}^K \omega_k(X) \frac{Y_i(1)^2}{e_k(X_i)} \right]$$

if $\exists(k, k') \in [K], e_k(X_i) \neq e_{k'}(X_i)$, and equality if $\forall k \in [K], e_k(X_i) = e(X_i)$, i.e. if the local propensity scores are all equal to one another. Then, with the same condition on strictness and equality,

$$\begin{aligned} \mathbb{E} \left[\frac{Y_i(1)^2}{e(X_i)} \right] &\leq \mathbb{E} \left[\sum_{k=1}^K \omega_k(X) \frac{Y_i(1)^2}{e_k(X_i)} \right] \\ &= \sum_{k=1}^K \mathbb{E} \left[\mathbb{E} \left[\mathbb{1}_{[H_i=k]} \mid X_i \right] \frac{Y_i(1)^2}{e_k(X_i)} \right] \\ &= \sum_{k=1}^K \mathbb{E} \left[\mathbb{E} \left[\mathbb{1}_{[H_i=k]} \frac{Y_i(1)^2}{e_k(X_i)} \mid X_i \right] \right] \\ &= \sum_{k=1}^K \mathbb{E} \left[\mathbb{1}_{[H_i=k]} \frac{Y_i(1)^2}{e_k(X_i)} \right] \\ &= \sum_{k=1}^K \rho_k \mathbb{E} \left[\frac{Y_i(1)^2}{e_k(X_i)} \mid H_i = k \right] \end{aligned}$$

Similarly, $\mathbb{E} \left[\frac{Y_i(0)^2}{1-e(X_i)} \right] \leq \sum_{k=1}^K \rho_k \mathbb{E} \left[\frac{Y_i(0)^2}{1-e_k(X_i)} \mid H_i = k \right]$. Then,

$$\begin{aligned} \mathbb{V} [\hat{\tau}^{\text{fed-IPW}^*}] &= \frac{1}{n} \left(\mathbb{E} \left[\frac{Y_i(1)^2}{e(X_i)} + \frac{Y_i(0)^2}{1-e(X_i)} \right] - \tau^2 \right) \\ &\leq \frac{1}{n} \left(\sum_{k=1}^K \rho_k \mathbb{E} \left[\frac{Y_i(1)^2}{e_k(X_i)} + \frac{Y_i(0)^2}{1-e_k(X_i)} \mid H_i = k \right] - \tau^2 \right) \\ &\leq \frac{1}{n} \left(\sum_{k=1}^K \rho_k \underbrace{\left(\mathbb{E} \left[\frac{Y_i(1)^2}{e_k(X_i)} + \frac{Y_i(0)^2}{1-e_k(X_i)} \mid H_i = k \right] - \tau_k^2 \right)}_{:=V_k} + \underbrace{\sum_{k=1}^K \rho_k \tau_k^2 - \tau^2}_{\mathbb{V}(\tau_H)} \right) \end{aligned}$$

On the other hand, by Lemma 2,

$$\mathbb{V} [\hat{\tau}^{\text{meta-IPW}^*}] = \frac{1}{n} \sum_{k=1}^K \rho_k V_k + \frac{1}{n} \mathbb{V} [\tau_H].$$

Hence $\mathbb{V} [\hat{\tau}^{\text{fed-IPW}^*}] = \mathbb{V} [\hat{\tau}^{\text{meta-IPW}^*}]$ if $\forall k \in [K], e_k = e$, and $\mathbb{V} [\hat{\tau}^{\text{fed-IPW}^*}] < \mathbb{V} [\hat{\tau}^{\text{meta-IPW}^*}]$ if $\exists(k, k') \in [K]^2, e_k \neq e_{k'}$.

AIPW. Similarly to the IPW case, we have

$$\begin{aligned} \mathbb{E} \left[\left(\frac{W_i(Y_i - \mu_1)}{e(X_i)} \right)^2 \right] &\leq \sum_{k=1}^K \rho_k \mathbb{E} \left[\left(\frac{W_i(Y_i - \mu_1)}{e_k(X_i)} \right)^2 \mid H_i = k \right], \\ \mathbb{E} \left[\left(\frac{(1 - W_i)(Y_i - \mu_0)}{1 - e(X_i)} \right)^2 \right] &\leq \sum_{k=1}^K \rho_k \mathbb{E} \left[\left(\frac{(1 - W_i)(Y_i - \mu_0)}{1 - e_k(X_i)} \right)^2 \mid H_i = k \right], \end{aligned}$$

and with Lemma 1:

$$\mathbb{V} [\tau(X_i)] = \sum_{k=1}^K \rho_k \mathbb{V} [\tau(X_i) \mid H_i = k] + \mathbb{V} [\tau_H].$$

Then, using Lemma 2, we have the desired result. $\square$

A.7. Proof of Theorem 5

Proof. Let $f : t \mapsto \frac{1}{t(1-t)}$ with $t \in]0, 1[$. f is convex. Then by Jensen's inequality with summing-to-one weights $\omega_k(X) = \mathbb{P}(H_i = k \mid X_i = X)$,

$$\begin{aligned} \mathcal{O}_{\text{global}} &= \mathbb{E} \left[f \left(\sum_{k=1}^K \omega_k(X) e_k(X) \right) \right] \\ &\leq \mathbb{E} \left[\sum_{k=1}^K \omega_k(X) f(e_k(X)) \right] \\ &\leq \sum_{k=1}^K \mathbb{E} [\mathbb{E} [\mathbb{1}_{[H_i=k|X_i]}] f(e_k(X))] \\ &\leq \sum_{k=1}^K \mathbb{E} [\mathbb{E} [\mathbb{1}_{[H_i=k]} f(e_k(X)) \mid X_i]] \\ &\leq \sum_{k=1}^K \mathbb{E} [\mathbb{1}_{[H_i=k]} f(e_k(X))] \\ &\leq \sum_{k=1}^K \mathbb{P}(H_i = k) \mathbb{E} [f(e_k(X)) \mid H_i = k] \\ &\leq \sum_{k=1}^K \rho_k \mathcal{O}_k \end{aligned}$$

$\square$

A.8. Proof of Theorem 6

Proof. By the Neyman orthogonality of the doubly robust AIPW influence function, the difference between the empirical estimator $\hat{\tau}_{\text{AIPW}}^{\text{fed}}$ and the oracle estimator $\hat{\tau}_{\text{AIPW}}^{\text{fed}*}$ is driven by the product of the estimation errors of the nuisance functions. Under cross-fitting, this yields the bound:

$$\sqrt{n} (\hat{\tau}_{\text{AIPW}}^{\text{fed}} - \hat{\tau}_{\text{AIPW}}^{\text{fed}*}) = \mathcal{O}_P \left(\sqrt{n} \|\hat{e} - e\|_2 \cdot \max_{w \in \{0,1\}} \|\hat{\mu}_w - \mu_w\|_2 \right).$$

To evaluate the L_2 error of the global propensity score, we decompose $\hat{e}(X) - e(X)$ by adding and subtracting $\sum_{k=1}^K \hat{\omega}_k(X) e_k(X)$ and apply the triangle inequality:

$$\begin{aligned} \|\hat{e} - e\|_2 &= \left\| \sum_{k=1}^K \hat{\omega}_k \hat{e}_k - \sum_{k=1}^K \omega_k e_k \right\|_2 \\ &\leq \sum_{k=1}^K \|\hat{\omega}_k (\hat{e}_k - e_k)\|_2 + \sum_{k=1}^K \|e_k (\hat{\omega}_k - \omega_k)\|_2 \\ &\leq \sum_{k=1}^K \|\hat{\omega}_k\|_\infty \|\hat{e}_k - e_k\|_2 + \sum_{k=1}^K \|e_k\|_\infty \|\hat{\omega}_k - \omega_k\|_2. \end{aligned}$$

Because the membership weights and local propensity scores are probabilities, they are bounded by 1 almost everywhere ($\|\hat{\omega}_k\|_\infty \leq 1$ and $\|e_k\|_\infty \leq 1$), yielding

$$\|\hat{e} - e\|_2 \leq \sum_{k=1}^K \|\hat{e}_k - e_k\|_2 + \sum_{k=1}^K \|\hat{\omega}_k - \omega_k\|_2.$$

Since $\liminf \frac{n_k}{n} > 0$, the local sample sizes grow proportionally to the global sample size n , guaranteeing that local estimators trained on n_k samples can achieve convergence relative to the global rate.

By assumption, the local nuisance estimators converge in L_2 norm at the $o_P(n^{-1/4})$ rate. Because K is finite, the global propensity score error preserves this rate, ensuring $\|\hat{e} - e\|_2 = o_P(n^{-1/4})$.

Since the global outcome regressions also converge at $\|\hat{\mu}_w - \mu_w\|_2 = o_P(n^{-1/4})$, the product of the errors is $o_P(n^{-1/2})$. The remainder term is thus $\mathcal{O}_P(\sqrt{n} \cdot o_P(n^{-1/2})) = o_P(1)$.

We conclude that $\sqrt{n}(\hat{\tau}_{\text{AIPW}}^{\text{fed}} - \tau) = \sqrt{n}(\hat{\tau}_{\text{AIPW}}^{\text{fed}*} - \tau) + o_P(1)$. By Slutsky's theorem, the empirical estimator converges to the exact same asymptotic distribution as the oracle estimator, achieving local semiparametric efficiency. $\square$

B. Heteroscedastic site effects

In this section, we consider a weaker version of Assumption 2.

Assumption 6 (Mean exchangeability). *For all $w \in \{0, 1\}$,*

$$\mathbb{E}\{Y(w) \mid X, H\} = \mathbb{E}\{Y(w) \mid X\}.$$

Unlike Assumption 2, this assumption does not require the full conditional distribution of $Y(w)$ to be invariant across sites. Higher-order moments may therefore vary with H . In particular, the conditional second moments may differ across sites,

$$\mathbb{E}\{Y(w)^2 \mid X, H = k\} \neq \mathbb{E}\{Y(w)^2 \mid X, H = k'\},$$

or equivalently the conditional variances may be site-dependent. This leads to heteroscedastic site effects.

This setting is relevant in applications where different sites measure outcomes with different levels of precision. For instance, hospitals may use measurement devices with distinct calibrations or noise profiles, inducing site-specific outcome variability even when the conditional mean outcome is stable across sites.

Under mean exchangeability, identification of the ATE and consistency of the proposed estimators are preserved. However, the oracle variances of the pooled estimators differ from those obtained under full distributional exchangeability:

Theorem 7. *Under Assumptions 1, 3 and 6, we have $\sqrt{n}(\hat{\tau} - \tau) \rightarrow \mathcal{N}(0, V^*)$ with*

$$\begin{cases} V_{\text{IPW}}^* = \mathbb{E} \left[\frac{e_H(X)Y(1)^2}{e(X)^2} \right] + \mathbb{E} \left[\frac{(1-e_H(X))Y(0)^2}{(1-e(X))^2} \right] - \tau^2, \\ V_{\text{AIPW}}^* = \mathbb{V}[\tau(X)] + \mathbb{E} \left[\frac{e_H(X)(Y(1)-\mu_1(X))^2}{e(X)^2} \right] \\ \quad + \mathbb{E} \left[\frac{(1-e_H(X))(Y(0)-\mu_0(X))^2}{(1-e(X))^2} \right]. \end{cases}$$

Under Assumption 6, the empirical federated AIPW estimator remains consistent and asymptotically normal. The identification argument is unchanged, and Neyman orthogonality still controls the effect of nuisance estimation errors under the same cross-fitted L_2 convergence conditions. Hence,

$$\sqrt{n}(\hat{\tau}_{\text{AIPW}}^{\text{fed}} - \tau) \rightsquigarrow \mathcal{N}(0, V_{\text{AIPW}}^*),$$

where V_{AIPW}^* denotes the oracle variance under site-level heteroscedasticity.

The usual empirical influence-function variance,

$$\hat{V}_{\text{AIPW}} = \frac{1}{n} \sum_{i=1}^n \left(\hat{\mu}_1(X_i) - \hat{\mu}_0(X_i) + \frac{W_i(Y_i - \hat{\mu}_1(X_i))}{\hat{e}(X_i)} - \frac{(1 - W_i)(Y_i - \hat{\mu}_0(X_i))}{1 - \hat{e}(X_i)} - \hat{\tau}_{\text{AIPW}}^{\text{fed}} \right)^2$$

is a heteroscedasticity-robust, sandwich-type estimator of the variance and does not require algorithmic modifications. It therefore accounts for site-specific outcome variability without explicitly modeling local second moments.

Proof. 1. Variance of the Oracle Fed-IPW Estimator

By definition, the asymptotic variance of the centralized oracle IPW estimator requires evaluating the second moment of its

summand. Using the Law of Iterated Expectations, we first condition on (X_i, H_i) , and apply SUTVA and Assumption 1:

$$\begin{aligned}
 \mathbb{E} \left[\left(\frac{W_i Y_i}{e(X_i)} \right)^2 \right] &= \mathbb{E}_{X_i, H_i} \left[\mathbb{E} \left[\left(\frac{W_i Y_i(1)}{e(X_i)} \right)^2 \mid X_i, H_i \right] \right] \\
 &= \mathbb{E}_{X_i, H_i} \left[\frac{\mathbb{E}[W_i \mid X_i, H_i] \mathbb{E}[Y_i(1)^2 \mid X_i, H_i]}{e(X_i)^2} \right] \\
 &= \mathbb{E}_{X_i, H_i} \left[\frac{e_{H_i}(X_i) \mathbb{E}[Y_i(1)^2 \mid X_i, H_i]}{e(X_i)^2} \right] \\
 &= \mathbb{E} \left[\frac{e_{H_i}(X_i) Y_i(1)^2}{e(X_i)^2} \right] \quad \text{Law of Iterated Expectations}
 \end{aligned}$$

Applying symmetric logic to the control arm and subtracting the squared mean, the variance of the oracle centralized multi-site IPW estimator is:

$$\mathbb{V} \left[\hat{\tau}^{\text{IPW}^*} \right] = \frac{1}{n} \left(\mathbb{E} \left[\frac{e_{H_i}(X_i) Y_i(1)^2}{e(X_i)^2} \right] + \mathbb{E} \left[\frac{(1 - e_{H_i}(X_i)) Y_i(0)^2}{(1 - e(X_i))^2} \right] - \tau^2 \right).$$

Multiplying by n yields the asymptotic variance V_{IPW}^* .

2. Variance of the Oracle Fed-AIPW Estimator

For the variance of the oracle multi-site centralized AIPW, we first note that the cross-covariance term A between the CATE and the augmentation terms is zero. Because Assumption 2 establishes mean exchangeability, we have $\mathbb{E}[Y_i(w) - \mu_w(X_i) \mid X_i, H_i] = 0$:

$$\begin{aligned}
 A &= \text{Cov} \left(\tau(X_i), \frac{W_i(Y_i - \mu_1(X_i))}{e(X_i)} - \frac{(1 - W_i)(Y_i - \mu_0(X_i))}{1 - e(X_i)} \right) \\
 &= \mathbb{E} \left[\tau(X_i) \left(\frac{W_i(Y_i(1) - \mu_1(X_i))}{e(X_i)} - \frac{(1 - W_i)(Y_i(0) - \mu_0(X_i))}{1 - e(X_i)} \right) \right] \\
 &= \mathbb{E} \left[\tau(X_i) \mathbb{E} \left[\frac{W_i(Y_i(1) - \mu_1(X_i))}{e(X_i)} \mid X_i, H_i \right] \right] - \mathbb{E} \left[\tau(X_i) \mathbb{E} \left[\frac{(1 - W_i)(Y_i(0) - \mu_0(X_i))}{1 - e(X_i)} \mid X_i, H_i \right] \right] \\
 &= \mathbb{E} \left[\tau(X_i) \frac{e_{H_i}(X_i) \mathbb{E}[Y_i(1) - \mu_1(X_i) \mid X_i, H_i]}{e(X_i)} \right] - \mathbb{E} \left[\tau(X_i) \frac{(1 - e_{H_i}(X_i)) \mathbb{E}[Y_i(0) - \mu_0(X_i) \mid X_i, H_i]}{1 - e(X_i)} \right] \\
 &= 0 - 0 = 0.
 \end{aligned}$$

Moreover, the variance of the treated augmentation term simplifies via the Law of Iterated Expectations, allowing us to collapse the expectation without assuming equivalent second moments across sites:

$$\begin{aligned}
 \mathbb{V} \left[\frac{W_i(Y_i(1) - \mu_1(X_i))}{e(X_i)} \right] &= \mathbb{E} \left[\frac{W_i(Y_i(1) - \mu_1(X_i))^2}{e(X_i)^2} \right] \\
 &= \mathbb{E}_{X_i, H_i} \left[\frac{\mathbb{E}[W_i \mid X_i, H_i] \mathbb{E}[(Y_i(1) - \mu_1(X_i))^2 \mid X_i, H_i]}{e(X_i)^2} \right] \\
 &= \mathbb{E}_{X_i, H_i} \left[\frac{e_{H_i}(X_i) \mathbb{E}[(Y_i(1) - \mu_1(X_i))^2 \mid X_i, H_i]}{e(X_i)^2} \right] \\
 &= \mathbb{E} \left[\frac{e_{H_i}(X_i) (Y_i(1) - \mu_1(X_i))^2}{e(X_i)^2} \right].
 \end{aligned}$$

Similarly, for the control term:

$$\mathbb{V} \left[\frac{(1 - W_i)(Y_i(0) - \mu_0(X_i))}{1 - e(X_i)} \right] = \mathbb{E} \left[\frac{1 - e_{H_i}(X_i)}{(1 - e(X_i))^2} (Y_i(0) - \mu_0(X_i))^2 \right].$$

Finally, leveraging independence across observations i and $A = 0$, the variance of the overall estimator is:

$$\begin{aligned}\mathbb{V}[\hat{\tau}^{\text{AIPW}*}] &= \mathbb{V}\left[\frac{1}{n}\sum_{i=1}^n\left(\mu_1(X_i) - \mu_0(X_i) + \frac{W_i(Y_i - \mu_1(X_i))}{e(X_i)} - \frac{(1 - W_i)(Y_i - \mu_0(X_i))}{1 - e(X_i)}\right)\right] \\ &= \frac{1}{n^2}\sum_{i=1}^n\left(\mathbb{V}[\tau(X_i)] + \mathbb{V}\left[\frac{W_i(Y_i - \mu_1(X_i))}{e(X_i)}\right] + \mathbb{V}\left[\frac{(1 - W_i)(Y_i - \mu_0(X_i))}{1 - e(X_i)}\right] + 2A\right) \\ &= \frac{1}{n}\left(\mathbb{V}[\tau(X_i)] + \mathbb{E}\left[\frac{e_{H_i}(X_i)}{e(X_i)^2}(Y_i(1) - \mu_1(X_i))^2\right] + \mathbb{E}\left[\frac{1 - e_{H_i}(X_i)}{(1 - e(X_i))^2}(Y_i(0) - \mu_0(X_i))^2\right]\right).\end{aligned}$$

Multiplying by n gives the stated asymptotic variance V_{AIPW}^* . $\square$

C. Federated learning of membership weights

We describe how to estimate membership weights using a general parametric classification model trained via Federated Averaging (FedAvg) (McMahan et al., 2017), without requiring access to individual-level data.

C.1. General Parametric Model

Let $\omega_k(X_i; \Theta)$ denote the predicted membership probability of site k for covariate vector X_i , where Θ are the parameters of the model. For a generic parametric model, $\Theta \in \mathbb{R}^p$ represents all trainable parameters, and the membership-weight estimation problem can be formulated as the minimization of the empirical cross-entropy loss:

$$\ell(\Theta; \mathcal{D}) = -\frac{1}{|\mathcal{D}|}\sum_{i=1}^{|\mathcal{D}|}\sum_{k=1}^K \mathbb{1}_{\{H_i=k\}} \log(\omega_k(X_i; \Theta)),$$

where H_i^{enc} is the one-hot encoding of the site membership H_i .

This framework encompasses a broad class of models, including:

- Multinomial logistic regression:

$$\omega_k(X_i; \Theta) = \frac{\exp(\theta_k^\top X_i)}{\sum_{k'=1}^K \exp(\theta_{k'}^\top X_i)}, \quad \Theta = (\theta_1, \dots, \theta_K) \in \mathbb{R}^{d \times K}.$$

- Neural networks, where Θ includes all weights and biases across layers, and $\omega_k(\cdot)$ is the softmax output of the final layer.

C.2. Federated Training Procedure

Algorithm 1 summarizes the generic Federated Averaging procedure for training any differentiable parametric model for membership-weight estimation.

Algorithm 1 Federated Learning of Membership Weights with a Parametric Model

Require: K sites, E local steps, η learning rate, T communication rounds, B batch size

```

1: Initialize global parameters  $\Theta_0$ 
2: for  $t = 1$  to  $T$  do
3:   for each client  $k \in [1, \dots, K]$  in parallel do
4:      $\Theta_{t+1}^{(k)} \leftarrow \text{LocalUpdate}(k, \Theta_t)$ 
5:   end for
6:    $\Theta_{t+1} \leftarrow \sum_{k=1}^K \frac{n_k}{n} \Theta_{t+1}^{(k)}$  // FedAvg aggregation
7: end for

8: Function  $\text{LocalUpdate}(k, \Theta)$ :
9:   for  $e = 1$  to  $E$  do
10:    Sample mini-batch  $\mathcal{B}_k$  of size  $B$  from local data  $\mathcal{D}_k$ 
11:    Compute predictions  $\omega_k(X_i; \Theta)$  for  $i \in \mathcal{B}_k$ 
12:    Compute gradient  $\nabla \ell(\Theta; \mathcal{B}_k)$  of cross-entropy loss
13:     $\Theta \leftarrow \Theta - \eta \nabla \ell(\Theta; \mathcal{B}_k)$ 
14:   end for
15: Return  $\Theta$ 
    
```

With a suitable choice of learning rate η and a moderate number of local steps E per round, Algorithm 1 converges to the same solution as centralized training as the number of communication rounds $T \rightarrow \infty$ (Stich, 2019; Khaled et al., 2020; Li et al., 2019). The resulting federated parameter estimate $\hat{\Theta}^{\text{fed}}$ yields the estimated membership weights:

$$\hat{\omega}_k(X_i) = \omega_k(X_i; \hat{\Theta}^{\text{fed}}), \quad k = 1, \dots, K.$$

Note that this federated procedure remains valid for learning $H|X$ despite each site k locally holding only one label ($H = k$).

C.3. Example: Multinomial Logistic Regression

When the model is multinomial logistic regression, Θ is the matrix of regression coefficients, and the local gradient takes the closed form:

$$\nabla \ell(\Theta; \mathcal{B}_k) = \frac{1}{B} \sum_{i \in \mathcal{B}_k} X_i (\omega(X_i; \Theta) - H_i^{\text{enc}}),$$

where $\omega(X_i; \Theta)$ is the vector of softmax probabilities for sample i .

D. Simulation Details

The parameters common to all settings are shown in Table 2, where

$$\gamma_2^{(\text{weak})} = [-2.5, -1, -0.15, -0.15, 0, -0.15, -1, -0.15, -0.15, 0]$$

and

$$\gamma_2^{(\text{good})} = [-.05, -.1, -.05, -.1, .05, -.1, -.05, -.1, .05, -.1]$$

.

DGP A-specific settings are shown in Table 3, where J_d is the $d \times d$ matrix of ones, and I_d is the $d \times d$ identity matrix.

DGP B-specific settings are shown in Table 4.

Real data. We classified a site as ‘large’ if it had at least $5 \times d$ observations in each treatment arm (minimum 85 per arm in our setting). For such sites, we estimated e_k with probability forests with 2,000 trees; otherwise, we used logistic regression. We estimated Fed-MW RF with a centralized random forest for $P(H|X)$ with 500 trees.

Parameter	Center 1	Center 2	Center 3
d	10		
$\mu_1(X)$	$\sum_{j=1}^5 \frac{j}{10} X_j^2 + \sum_{j=6}^{10} \frac{j}{10} X_j + X_9 * X_{10}$		
$\mu_0(X)$	$\sum_{j=1}^5 \frac{3j-10}{30} X_j^2 + \sum_{j=6}^{10} \frac{3j-10}{30} X_j + X_1 * X_{10}$		
e_k	$\text{Logistic}(x, \gamma_k)$		
γ_k	$[-.25, .25, -.25, -.25, .25, -.25, -.25, .25, -.25, .25]$	No overlap: not logistic (only controls) Weak overlap: $\gamma_2^{(\text{weak})}$ Good overlap: $\gamma_2^{(\text{good})}$	$[.15, -.15, .15, -.15, .15, -.15, .15, -.15, .15, -.15]$


Table 2. Common simulation parameters.

Parameter	Center 1	Center 2	Center 3
n_k	650		
$\mathcal{D}_k$	$\mathcal{N}(\mu_k, \Sigma_k)$		
μ_k	$(1, \dots, 1) \in \mathbb{R}^d$	$(1.5, 1.5, 1.5, 1, \dots, 1) \in \mathbb{R}^d$	$(2, 2, 2, 1, \dots, 1) \in \mathbb{R}^d$
Σ_k	$I_d + 0.5J_d$	$0.6I_d + 0.4J_d$	$3I_d + 0.3J_d$

Table 3. Simulation parameters specific to DGP A.


E. Overlap Improvement

Echoing theorem 5, Figures 5 and 6 display the empirical distributions (on a log-scale) of the local propensity score in site 2 (e_2) and of the global propensity score e in the *Poor local overlap* scenario (see Figure 2b for corresponding results). A good overlap is when both the propensity score distributions for the treated and control overlap, and are far from 0. We see that the poor overlap at site 2 (Figure 5), with values of e_2 on the local data close to 0, is significantly improved at the global level (Figure 6).


 Figure 5. Local overlap in site 2 for the *Poor local overlap* scenario (DGP A).

Parameter	Value
n	4000
$\mathcal{D}$	$\frac{2}{3}\mathcal{N}(\mu_1, \Sigma_1) + \frac{1}{3}\mathcal{N}(\mu_2, \Sigma_2)$
μ_1	$(0, \dots, 0) \in \mathbb{R}^d$
μ_2	$(1.5, \dots, 1.5) \in \mathbb{R}^d$
Σ_1	I_d
Σ_2	$I_d + 0.5J_d$
$\mathbb{P}(H_i = k \mid X)$	$\text{Logistic}(x, \theta_k)$
θ_1	$[-0.5, -0.5, 0.2, -0.5, -0.5, 0.2, -0.5, -0.5, 0.2, 0.2]$
θ_2	$[0.5, 0.5, 0.2, 0.5, 0.5, 0.2, 0.5, 0.5, 0.2, 0.5]$
θ_3	$[1, 1, 0.2, 0.2, 0.2, 0.2, 0.2, 0.2, 0.2, 0.2]$

Table 4. Simulation parameters specific to DGP B.


 Figure 6. Global overlap for the *Poor local overlap* scenario (DGP A). We see a clear improvement compared to the local overlap in site 2 (Figure 5).

F. Additional Simulations

F.1. Large number of sites

We further investigate the robustness of our estimators as the number of sites increases. Specifically, we simulate $K = 20$ sites under weak local overlap (corresponding to the setting of Figure 2b in the main paper) and report the bias and mean squared error (MSE) for both data-generating processes (DGP A and DGP B).

Table 5. Bias and mean squared error (MSE) under $K = 20$ sites with weak local overlap. Bold values indicate unbiasedness or lowest MSE.

Method	Bias (DGP A)	MSE (DGP A)	Bias (DGP B)	MSE (DGP B)
Oracle IPW	0.00	0.298	0.00	0.130
Fed. IPW-MW (logistic)	0.48	0.294	0.00	0.095
Fed. IPW-MW (NN)	0.00	0.240	0.00	0.123
Fed. IPW-DW	0.00	0.096	0.50	0.450
Meta-SW IPW	-2.09	13.279	-1.77	7.977
Oracle AIPW	0.00	0.007	0.00	0.049
Fed. AIPW-MW (logistic)	0.79	0.639	0.00	0.078
Fed. AIPW-MW (NN)	0.00	0.153	0.00	0.084
Fed. AIPW-DW	0.00	0.019	0.25	0.194
Meta-SW AIPW	-1.24	1.775	-2.36	18.465

As shown in Table 5, estimators using neural network membership weights consistently achieve unbiasedness across both DGPs, although with slightly higher variance compared to their well-specified parametric counterparts. This is consistent with the behavior of more flexible models, which trade off a small increase in variance for robustness to misspecification. When parametric assumptions are correctly specified (e.g., Gaussian density weights in DGP A, logistic membership weights in DGP B), these methods achieve the lowest MSE. Overall, the NN-based approach demonstrates robust consistency across DGPs and scales well with the number of sites.

F.2. Fully model-agnostic ATE estimation

We additionally investigate the benefits of fully model-agnostic estimation for membership weights $P(H | X)$. In this setting, we use federated neural networks to estimate membership probabilities in a flexible way and pair them with non-parametric local propensity score estimators (e.g., Random Forests). This procedure does not require prior parametric knowledge about $H | X$ or $X | H$, making it robust to a wide range of DGPs.

We simulate $K = 3$ sites with $d = 10$ covariates, $n_k = 600$ observations per site, and $X | H = k \sim D_k$ as a mixture of Gaussians with weak local overlap. The results are summarized in Table 6.

Method	Bias	MSE
Oracle IPW	0.00	0.116
Fed. IPW-MW (NN MW + RF e_k)	0.00	0.075
Meta-SW IPW (RF e_k)	0.89	0.884
Oracle AIPW	0.00	0.030
Fed. AIPW-MW (NN MW + RF e_k)	0.00	0.040
Meta-SW AIPW (RF e_k)	-0.36	0.180

Table 6. Bias and MSE under fully non-parametric estimation of membership weights and local propensity scores ($K = 3$, weak local overlap). Bold values indicate unbiased (or nearly unbiased) estimators.

As shown in Table 6, our federated approach with neural-network-based membership weights and random-forest local propensities yields nearly unbiased estimates with substantially lower MSE compared to meta-analysis estimators. This demonstrates that misspecifications in either the membership weights or the local propensity scores can be effectively mitigated by adopting fully non-parametric methods in a federated learning setting.

F.3. Robustness to local misspecifications

Similar to meta-analysis estimators, we argue that inconsistencies in local propensity models can be mitigated as K grows, provided that such inconsistencies are marginal relative to the number of correctly specified sites. To illustrate this, we simulate K sites under DGP B, each holding $n_k = 60$ observations drawn from homogeneous covariate distributions and




Figure 7. Graphical model for multi-site observational data with site covariates Z mediating the site effect.

sharing the same true propensity model (logistic). We then intentionally misspecify the propensity model for site $k = 1$, setting $\hat{e}_1(x) \equiv 0.1$ for all x .

Table 7 reports the bias and mean squared error (MSE) of several estimators under two scenarios: $K = 3$ (few sites) and $K = 200$ (many sites).

Method	Bias ($K = 3$)	MSE ($K = 3$)	Bias ($K = 200$)	MSE ($K = 200$)
Oracle IPW	0.00	0.272	0.00	0.002
Fed. IPW (logistic MW)	0.38	0.344	0.00	0.002
Meta IPW	2.61	7.902	0.03	0.003
Oracle AIPW	0.00	0.058	0.00	0.001
Fed. AIPW (logistic MW)	-0.33	0.325	-0.01	0.002
Meta AIPW	-0.51	0.481	-0.07	0.007

Table 7. Bias and mean squared error (MSE) under local misspecification of site $k = 1$ propensity model.

With only $K = 3$ sites, both federated and meta-analysis estimators exhibit substantial bias. However, as the number of sites increases, the influence of the single misspecified site becomes negligible, and the bias decreases markedly (e.g., from 0.38 to 0.00 for Federated IPW with logistic membership weights). As expected, the AIPW estimator retains its double-robustness property even under misspecification of the outcome model, leading to consistently low bias and MSE across settings.

G. Additional Details on Handling Site Effects

G.1. Federated Tests for Site Effects

Assumption 2 posits that site membership H provides no information about potential outcomes conditional on covariates X and treatment W . This assumption is *testable* by fitting a federated regression model of Y against (X, H) , and testing whether the coefficients for H are jointly zero. In federated learning, this test is implemented via one-shot sharing of sufficient statistics: each site transmits its local covariance matrix $\hat{\Sigma}_k = \frac{1}{n_k} \sum_{i=1}^{n_k} X_i^\top X_i$ and vector $\hat{\gamma}_k = \frac{1}{n_k} \sum_{i=1}^{n_k} X_i^\top Y_i$. The server aggregates these to compute a joint F-test or Wald test on the H block, yielding a p -value. We recommend adapting the estimation strategy based on this test’s result.

G.2. Experimental Details

We provide the simulation parameters used in Table 1. We simulate $K = 30$ sites, each with $n_k \in \{20, 60, 90\}$ observations. Covariates $X \in \mathbb{R}^3 \sim \mathcal{N}(0, I_3)$ are drawn i.i.d. and allocated to sites via $H \sim \text{Logit}(\theta^\top X)$ to induce covariate shift. Site-level covariates $Z \in \mathbb{R}^3$ (site effects) are drawn from site-specific Gaussians: for each site k , a centroid $\mu_k \sim \mathcal{N}([1, 1, 1], I_3)$ is drawn, then $Z \mid H = k \sim \mathcal{N}(\mu_k, I_3)$. Treatment is assigned by site-specific logistic models. Outcomes are linear: $Y = X^\top \beta + Z^\top \delta + \tau W + \varepsilon$, with $\varepsilon \sim \mathcal{N}(0, 1)$, $\tau = 10$, and $\beta = \delta = [-2, -1.33, -0.66]$.

We estimate membership probabilities $P(H \mid \cdot)$ using multinomial logistic regression or a two-layer neural network (NN), trained on either (X) or (X, Z) . Local propensities $\hat{e}_k(X)$ are fitted via logistic regression, and federated outcome models




Figure 8. Graphical model for multi-site observational data where site effects are not fully explained by Z (direct $H \rightarrow Y$ edge).

$\hat{\mu}_w(X, Z)$ via OLS. Membership weights are aggregated to form $\hat{e}(x, z) = \sum_k \hat{P}(H=k \mid x, z) \hat{e}_k(x)$. IPW and AIPW estimators are computed as $\hat{\tau}_{\text{IPW}} = \frac{1}{n} \sum_i \frac{W_i Y_i}{\hat{e}(X_i)} - \frac{1}{n} \sum_i \frac{(1-W_i) Y_i}{1-\hat{e}(X_i)}$ and $\hat{\tau}_{\text{AIPW}} = \frac{1}{n} \sum_i \left[\frac{W_i(Y_i - \hat{\mu}_1)}{\hat{e}(X_i)} - \frac{(1-W_i)(Y_i - \hat{\mu}_0)}{1-\hat{e}(X_i)} + (\hat{\mu}_1 - \hat{\mu}_0) \right]$. Meta-analysis estimators adjust only for X , as averaging over H inherently stratifies by site.

G.3. On Unexplained Site Effects

If site-level covariates Z are unavailable or insufficient to explain all site effects (e.g., due to a direct $H \rightarrow Y$ edge, see Figure 8), the core identification assumption remains Assumption 1: $\{Y(1), Y(0)\} \perp W \mid (X, H)$. In this regime, nuisance functions—propensity scores $e(X, Z, H)$ and outcome models $\mu_w(X, Z, H)$ —must be learned federatively using parametric models that explicitly include site fixed effects (H). Including Z in these models helps to reduce the variance of the nuisance functions. Note that the theoretical and practical advantage of federated learning over meta-analysis is diminished in this setting, as meta estimation effectively stratifies inference on H .