HERMES: Heterogeneous Effects Representation with Matched Embeddings using Siamese Networks

Anonymous authors
Paper under double-blind review

Abstract

A growing literature has focused on representation learning as it relates to the estimation and analysis of heterogeneous treatment effects in both experimental and observational settings. We are specifically interested in the estimation of conditional average treatment effects (CATE) functions, i.e. functions mapping the effect of a binary treatment to the space of unit-level covariates. In the absence of a controlled randomized mechanism of treatment assignment, simple comparisons between treated and control populations can be potentially confounded by significant distributional differences in the covariate space. In this context, recent *representation learning* strategies aim to learn balanced latent representations in a new space where the treated and control distributions are more comparable, addressing confounding and reducing global bias. In this work, we leverage self-supervision and contrastive learning and propose a novel contrastive loss function that structures the latent space according to the similarity of estimated *individual treatment effects* (ITE). We integrate this contrastive learning approach in HERMES (*Heterogeneous Effects Representation with Matched Embeddings using Siamese networks*), a Siamese Neural Network which learns a structured latent space by dynamically pairing samples whose estimated ITEs are similar. Unlike representation learning approaches that rely only on covariates, HERMES injects the ITE into representation learning, improving accuracy under standard assumptions. Experiments on IHDP and JOBS benchmarks show that HERMES improves the expected Precision in MSE by 14–15% over baselines, without added inference cost.

1 Introduction

Estimating the heterogeneous causal effect of a binary intervention is a fundamental task in many policy and scientific domains. We are specifically interested in estimates of the conditional average treatment effect (CATE) as the main estimand of interest in personalized treatment regimes (Chen & Xie, 2023). In *precision medicine*, for example, targeted cancer therapies may be effective only for patients with certain mutations. Administering treatment indiscriminately exposes patients to unnecessary toxicity and risks. Estimating the CATE allows us to identify who is most likely to benefit from intervention (Harvey et al., 2020). Similar considerations apply whenever decision making is likely to ensue in heterogeneous unit-specific effects. See, for example, Athey & Imbens (2017) and LaLonde (1986) for applications of CATE to *policy-making and welfare*, or Ascarza (2018) for applications to *marketing strategies*.

Carefully designed randomized controlled trials (RCT) are the gold standard for measuring causal effects. However, these studies are often hindered by numerous practical limitations (Munk et al., 2020; Klonoff, 2020), and many inquiries must rely on the information included in *observational data*. Estimating CATE functions from observational data must, however, contend with a fundamental challenge: *confounding*. This arises when covariates that influence treatment assignment are also associated with the outcome, making it difficult to isolate the causal effect of treatment.

A common framework for careful causal analysis in observational settings relies on the idea of potential outcomes Imbens & Rubin (2015), where every observational unit is endowed with two "potential outcomes" corresponding to observations obtained with and without treatment. Clearly, each realized sample only

allows for the observation of one of these two outcomes. Therefore, the unobserved "*counterfactual*" outcome proscribes direct observation of the individual treatment effect (ITE). While this simple idea highlights why this problem differs from standard supervised learning, whenever treatment assignment and outcomes can be assumed conditionally independent, given measurement of a sufficient set of features (*unconfoundedness*), the causal effect is indeed identifiable through *local average differences*.

Crucially, local averaging aims to closely emulate the internal validity of a randomized controlled trial, and produce "locally" accurate estimates of the individual treatment effect despite the inaccessibility of the counterfactual outcome. To achieve this, the discrepancy in feature distributions for the treated and control subject groups are expressed via a balancing score (Rosembaum & Rubin, 1983), which is then used to apply a correction to the estimate of conditional outcome functions (Stuart, 2010; Rotnitzky & Robins, 1995; Athey et al., 2017).

Several methods trained to separately learn a balancing score and conditional outcome functions can be adapted to obtain consistent estimators (Künzel et al., 2019). In contrast, the representation learning literature has focused on end-to-end learning of CATE functions with balanced latent representations of the feature space (Shalit et al., 2017). Within this framework, the embedding distributions of treated and control individuals are as similar as possible, thus addressing the problem of confounding and reducing global bias. Learning representations in which treated and control individuals are approximately indistinguishable with respect to pre-treatment covariates enables valid comparison: when their latent distributions match, each treated individual effectively has a counterpart in control with equivalent features, mitigating confounding due to observed covariates.

This is often achieved using Neural Networks as function learners with balanced representations of the feature space learned through integral probability metric (IPM) regularization (Shalit et al., 2017), or by framing the problem in a multi-task learning context with shared embeddings predicting both treatment assignment and outcomes (Shi et al., 2019; Tesei et al., 2023). This strategy is justifiable from a theoretical perspective in that the expected error in learning ITE is upper bounded by two components: the error in learning both the factual and counterfactual outcomes plus a term quantifying the dissimilarity of the treated and control distributions of the learned embeddings (Shalit et al., 2017).

These representation strategies are, however, not designed to maintain local coherence for individuals with similar effects, and so, micro-regions exist in the latent space where two nearly identical individuals are projected far apart if their ITEs differ. Thus, we hypothesize the possibility to build and train a model on the correlation between latent space and ITE, in order to create representations whose similarity reflects a similar ITE value. In this work, we propose the estimation of heterogeneous treatment effects via representation learning with matched embeddings using Siamese networks (HERMES), whose goal is to learn a representation such that individuals with "similar" treatment effects are brought closer together in latent space, while individuals with very different effects are distant by at least an *adaptive margin*.

We define a *contrastive loss* based on: *(i) positive pairs*, i.e., individuals with similar ITE value, and *(ii) negative pairs*, i.e., individuals with ITE values difference over a threshold p. To build such pairs, no external labels are needed: covariates X and treatment T suffice, using the ITE values computed from the model's own predictions. The learning signal comes entirely from the data, without leveraging the observed outcome Y. Unlike classic representation learning approaches that rely only on covariates, `HERMES` injects treatment-effect signals into representation space, using a contrastive term which preserves overall moment-matching between treated and controls, while enforcing tight local alignment within positive pairs and margin-based separation for negative pairs.

On IHDP and JOBS benchmarks, `HERMES` outperforms `BCAUSS`, by leveraging local estimators that exploit the proximity of individuals in the latent space to produce more stable estimations.

2 Problem Setup and Related Work

We follow the potential outcomes framework of Rubin (1974) and consider a target population P from which a realization of n independent random variables $\{(X_i, T_i, Y_i(0), Y_i(1))\}_{i=1}^n \sim P$ is obtained as training data. Here, $X_i \in \mathbb{R}^d$ is a vector of d pre-treatment covariates, $T_i \in \{0, 1\}$ is the treatment received (0 for

control, 1 for treated), $Y_i(0) \in \mathbb{R}$ is the is the potential outcome for unit i when assigned to the control group, and $Y_i(1) \in \mathbb{R}$ is the potential outcome for unit i when assigned to the treatment group. We wish to make inference on the individual treatment effect (ITE) $D_i := Y_i(1) - Y_i(0)$, however, we only observe $Y_i = T_i\, Y_i(1) + (1 - T_i)\, Y_i(0)$, which highlights the main difference between ITE estimation and standard supervised learning.

Some inferential progress can be made by assuming *unconfoundedness*, also known as *strong ignorability of treatment assignment* (Rosenbaum & Rubin, 1983),

$$\{Y(0), Y(1)\} \perp\!\!\!\perp T \mid X, \text{ and } 0 < Pr(T = 1 \mid X) < 1.$$

When the conditional means $E(Y(0) \mid X = x)$ and $E\{Y(1) \mid X = x)$ exist, a more wieldy estimation target is the *conditional average treatment effect* (CATE),

$$\tau(x) = E\{Y(1) \mid X = x\} - E\{Y(0) \mid X = x\} = E\{Y(1) - Y(0) \mid X = x\},$$

which quantifies the expected treatment effect evaluated at a specific point in the feature space $x \in \mathcal{X}$. Crucially, the strong ignorability assumption is sufficient to identify the CATE (Künzel et al., 2019), and all that remains is to adapt statistical methods for the estimation of treatment effects to obtain CATE estimates with good asymptotic and finite sample properties.

2.1 Representation Learning for CATE Estimation

Access to sufficient covariate information is essential to remove confounding bias, allowing the statistical identification of CATE functions. In most applications, however, one only has access to a vector of potential confounders X, and it is often more efficient to forego conditioning on the entire information in X, and instead work with a reduced dimensional representation $\Phi(X)$. The goal is twofold: first, the representation $\Phi(X)$ must capture all the information relevant for predicting the outcome; second, it must remove the confounding bias. This latter objective translates into finding a balanced representation, whose distribution is the same for the treated and control groups.

Formally, a representation function $\Phi : \mathcal{X} \to \mathcal{R}$, where $\mathcal{R}$ is the representation space, is a twice differentiable, one-to-one function. The function Φ pushes forward the treated and control distributions into $\mathcal{R}$; we denote the induced distribution by p_Φ.

Definition 1 *We indicate with $p_\Phi^{T=1}(r) = p_\Phi(r \mid T = 1)$ and $p_\Phi^{T=0}(r) = p_\Phi(r \mid T = 0)$ the treated and control distributions induced by the representation Φ over the latent space $\mathcal{R}$.*

Intuitively, Φ maps the original covariates $\mathcal{X}$ into $\mathcal{R}$, where each individual is represented by a vector $r = \Phi(x)$. The distributions $p_\Phi^{T=1}(r)$ and $p_\Phi^{T=0}(r)$ describe how the treated and control individuals are distributed in $\mathcal{R}$.

An important class of representation functions are defined in relation to the concept of balancing score (Rosembaum & Rubin, 1983).

Definition 2 (Balancing score) *A function $b(X)$ is called* balancing score *if :* $X \perp\!\!\!\perp T \,\big|\, b(X)$.

This implies that, among individuals with the same value of $b(X)$, the distribution of covariates X is identical between treated and control. Once the value of $b(X)$ is known, the treatment assignment T provides no additional information about X. As a result, individuals with the same balancing score can be considered comparable, regardless of whether they received the treatment.

The best known case of balancing score is the *propensity score* $e(X) := P\big(T = 1 \mid X\big)$.

These types of representation function play a crucial role in CATE estimation. Specifically, reminding ourselves that the observable outcome is denoted with $Y = T\,Y(1) + (1 - T)\,Y(0)$, we restate the classical result of (Rosenbaum & Rubin, 1983).

Theorem 1 (Rosenbaum & Rubin, 1983) *If the information in X ensures strong ignorability of treatment assignment and $b(X)$ is a balancing score, then for each fixed value u of $b(X)$, we have:*

$$\underbrace{\mathbb{E}\{Y(1) \mid b(X){=}u\} - \mathbb{E}\{Y(0) \mid b(X){=}u\}}_{\textit{causal effect at score } u} = \underbrace{\mathbb{E}\{Y \mid b(X){=}u,\, T{=}1\} - \mathbb{E}\{Y \mid b(X){=}u,\, T{=}0\}}_{\textit{observed difference at score } u}.$$

The theorem states that within groups of individuals with the same balancing score, the difference in observed outcomes between treated and control individuals coincides with the causal effect of the treatment. Crucially, while the left-hand side is expressed in terms of a counterfactual quantity, the right-hand side is expressed as the difference between observable quantities, and therefore amenable to estimation.

Beyond CATE, a common downstream estimand is the *average treatment effect* (ATE), defined by averaging over $\mathcal{X}$, s.t.:

$$\psi \;=\; E[\,E\{Y \mid X,\, T{=}1\} - E\{Y \mid X,\, T{=}0\}] = E[\,E\{Y \mid b(X),\, T{=}1\} - E\{Y \mid b(X),\, T{=}0\}]\,.$$

Methods such as *back-door adjustment* (Belloni et al., 2013; Chernozhukov et al., 2024) and *weighting* approaches (Austin, 2011) use the propensity score to re-weight observational data, ensuring comparability between treated and control groups. *Targeted Learning* (van der Laan & Rose, 2018) refines initial estimators through a second-stage model to achieve an optimal bias–variance trade-off, while other approaches design treatment-effect–specific splitting criteria (Su et al., 2009; Zhang et al., 2017) or leverage ensemble/meta-algorithms (Künzel et al., 2019; Wager & Athey, 2017). Neural network–based solutions, such as *Propensity Dropout* (Alaa et al., 2017) and *Perfect Matching* Schwab et al. (2019), integrate propensity scores into their architectures. Here, we focus on *representation learning methods* through the formalism introduced by Shalit et al. (2017).

2.2 Representation Functions and Errors in Estimation

Let $\Phi(X)$ be a representation function and denote the conditional expected outcome with $h_t\{\Phi(x)\} := E\{Y \mid T = t, \Phi(X) = \Phi(x)\}$. If $\Phi(X)$ is a balancing score, and $\hat{h}_t : \mathcal{T} \times \mathcal{X} \to \mathcal{Y}$ is an estimator of the conditional outcome, we may consider CATE estimators of the form:

$$\hat{\tau}_\Phi(x) = \hat{h}_1\{\Phi(x)\} - \hat{h}_0\{\Phi(x)\}.$$

Crucially, even under balancing requirements, the representation function $\Phi(X)$ is not unique. We are, therefore, interested in learning $\Phi(X)$ to minimize the expected mean square error (EMSE):

$$\text{EMSE}_\Phi = E\left[\{\tau(x) - \hat{\tau}_\Phi(x)\}^2\right].$$

In this context, a fundamental framework for representation learning in causal inference was introduced by Shalit et al. (2017), who developed an upper bound for the EMSE, which can be expressed as the sum of factual losses and a measure of distributional imbalance between treated and control subpopulations.

Let $L : \mathcal{Y} \times \mathcal{Y} \to \mathbb{R}_+$ be a *loss function*. The expected factual treated and control losses are defined as:

$$\varepsilon_F\left\{\hat{h}_1(\Phi)\right\} \;:=\; E\left[L\left\{Y, \hat{h}_1(\Phi(X))\right\} \mid T = 1\right], \qquad \varepsilon_F\left\{\hat{h}_0(\Phi)\right\} \;=\; E\left[L\left\{Y, \hat{h}_0(\Phi(X))\right\} \mid T = 0\right].$$

The term $\varepsilon_F\{\hat{h}_1(\Phi)\}$ represents the expected predictive error restricted to the treated group. Analogously, $\varepsilon_F\{\hat{h}_0(\Phi)\}$ captures the expected predictive error restricted to the control group. Together, these two losses provide a group-specific decomposition of the overall factual risk.

Distributional imbalance may be quantified using *Integral Probability Metrics* (IPM). Specifically, for two probability density functions p, q defined over $S \subseteq \mathbb{R}^d$, and for a family $\mathcal{G}$ of functions $g : S \to \mathbb{R}$, we define the IPM between probability distributions as:

$$\text{IPM}_\mathcal{G}(p, q) \;=\; \sup_{g \in \mathcal{G}} \left| \int_S g(s)\big(p(s) - q(s)\big)\, ds \right|.$$

The choice of the function class $\mathcal{G}$ determines the specific form of the metric. For example, if $\mathcal{G}$ is the set of 1-Lipschitz functions, the IPM reduces to the Wasserstein distance; if $\mathcal{G}$ is the unit ball of a reproducing kernel Hilbert space (RKHS), the IPM reduces to the well-known *Maximum Mean Discrepancy* (MMD).

Below we restate a simplified version fo the result by Shalit et al. (2017), which will serve as a guiding principle for the design of CATE estimation strategies.

Theorem 2 *Let $\Phi : \mathcal{X} \to \mathcal{R}$ be a one-to-one representation function, σ_Y^2 be a residual variance parameter, and $\mathcal{G}$ be a suitable family of functions. Then, under squared error loss, there exists a constant $B_\Phi > 0$ s.t.:*

$$EMSE_\Phi \;\leq\; 2\Big[\varepsilon_F\left\{\hat{h}_1(\Phi)\right\} + \varepsilon_F\left\{\hat{h}_0(\Phi)\right\} + B_\Phi \cdot \mathrm{IPM}_{\mathcal{G}}\big(p_\Phi^{T=1},\, p_\Phi^{T=0}\big) - 2\sigma_Y^2\Big], \tag{1}$$

Theorem 2 establishes that minimizing the EMSE requires two complementary objectives: achieving low prediction error on the observed outcomes for both treated and control individuals, and ensuring that the learned representation aligns the distributions of treated and untreated populations in the representation space, as measured by the IPM.

2.3 CATE and Deep Neural Representation Learning

The full expressivity of deep learning strategies has proven to be of crucial importance in the estimation of the cate function $\tau(x)$. One of the earliest contributions in this domain is the `TARNet` structure introduced by Shalit et al. (2017) and its refined counterpart (`CFRNet`) (Shalit et al., 2017).

Both methods obtain stable estimates of the conditional outcome functions $h_1(\Phi)$ and $h_0(\Phi)$ by sharing a common representation of the covariates. A "Y-shaped" architecture is proposed; consisting of a shared encoder $\Phi(X)$ that feeds two separate potential-outcome heads, h_0 and h_1. The encoder maps the input covariates X to a latent representation $\Phi(X)$. The two heads then take this representation to estimate the potential outcomes: $\hat{Y}(1) = h_1(\Phi(X))$ and $\hat{Y}(0) = h_0(\Phi(X))$.

`TARNet` is trained end-to-end to minimize the regression error for *factual* (observed) outcomes. While deploying the same architecture, `CFRNet` encourages balanced representations by introducing a loss function, which adds a regularization term to the factual loss and explicitly penalizes the discrepancy between distributions, measured as:

$$\mathcal{L}_{\mathrm{CFR}} = \mathcal{L}_{\mathrm{factual}} + \alpha \cdot \mathrm{MMD}^2\{\Phi(\mathbf{X}_{T=1}), \Phi(\mathbf{X}_{T=0})\}, \tag{2}$$

where $\mathcal{L}_{\mathrm{factual}}$ is the loss function of `TARNet`, MMD is the *Maximum Mean Discrepancy*, and α is a hyperparameter controlling the strength of the balancing constraint. By forcing the latent distributions to overlap, `CFRNet` mitigates the out-of-distribution problem that plagues `TARNet`. However, this global alignment may prove to be an ineffective, since MMD operates at an aggregate level, ensuring the overall distributions match, but it provides no guarantee of *local* balance.

While `CFRNet` attempts to solve confounding by forcing balance, others have proposed deep learning architectures aimed explicitly at the construction of representation functions with balancing score properties. `DRAGONNET` (Shi et al., 2019), for example, extends the "Y-shaped" architecture into a "trident" by adding a third head dedicated to estimating a propensity score. In this framework, training is based on loss functions aimed at ensuring asymptotic efficiency and double rubustness through targeted regularization (van der Laan & Rose, 2018). A similar approach was adopted in Tesei et al. (2023) - `BCAUSS`, which uses the architecture of `DRAGONNET` but learns a *balancing score* $g(\mathbf{x})$ through a loss function $\mathcal{L}_{\mathrm{BAL}}$ which minimizes the squared difference between the weighted means of the covariates in treated (t_i) and control ($1 - t_i$) groups:

$$\mathcal{L}_{\mathrm{BAL}} = \frac{1}{d}\sum_{k=1}^{d}\left(\frac{\sum_{i=1}^{n}\frac{t_i}{g(\boldsymbol{x}_i)}x_{i,k}}{\sum_{i=1}^{n}\frac{t_i}{g(\boldsymbol{x}_i)}} - \frac{\sum_{i=1}^{n}\frac{1-t_i}{1-g(\boldsymbol{x}_i)}x_{i,k}}{\sum_{i=1}^{n}\frac{1-t_i}{1-g(\boldsymbol{x}_i)}}\right)^2. \tag{3}$$

Here, direct optimization in the balance of pretreatment covariates proves to enhance robustness in empirical scenarios with poor distributional overlap (Fan et al., 2023).

The idea of overlapping (positive) vs. non-overlapping (negative) samples has also been exploited using contrastive learning approachs, e.g. in the *Contrastive Individual Treatment Effects* (`CITE`) method proposed by

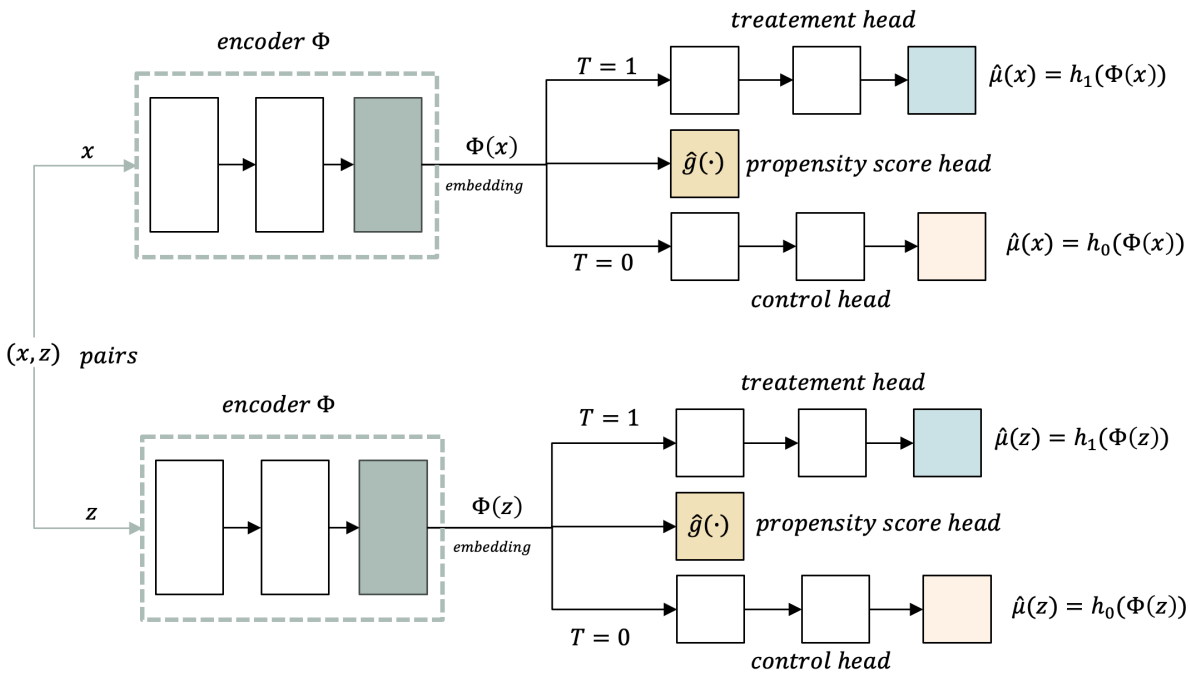

Figure 1: Overview of HERMES: a pair (x, z) of covariate vectors is passed through two branches; each branch consists of an *encoder* $\Phi(\cdot)$ that yields the latent embedding $\Phi(x)$, two *treatment-specific heads* to estimate the potential outcomes $\hat{\mu}_T(x) = h_T(\Phi(x))$, and a third head which calculates the propensity score $\hat{g}(\cdot)$; the output of the network is a ITE-similarity score between x and z.

Li & Yao (2022). Its core idea is to use the propensity score to define *positive* (individuals with propensity scores near 0.5, i.e., the "ideal" balanced case) and *negative* (individuals with scores near 0 or 1. i.e., the sources of bias) examples. Using such pairs, a noise contrastive loss forces the model to push the representations of positive pairs closer together (attraction) while simultaneously pushing apart the representations of negative pairs (repulsion) in the latent space. While potentially appealing, this approach is entirely dependent on an external, auxiliary, pre-trained propensity score model, which may influence the quality of the final learned representation.

Within this contextual setting, our work aims to address the central challenge posed by CATE estimation: *unreliable extrapolation*. This occurs when a model must predict a treatment effect for an individual without having similar "counterfactual twins" in the comparison group, forcing it to make risky assumptions based on general trends. In the following sections, we introduce HERMES, a framework that combines self supervised and contrastive learning, for globally and locally balanced representations of covariate information, leading to robust and improved inference for CATE functions.

3 CATE Representation Learning via Matched Siamese Embeddings

Theorem 2 shows that reducing the EMSE requires *both (i)* low factual prediction errors and *(ii)* a small discrepancy between the latent distributions of treated and controls. Our hypothesis is that we can further reduce representational discrepancies by introducing a *contrastive loss* function. The intuition is to force the model to create a latent space $\mathcal{R}$ in which the notion of "distance in representations" is semantically linked to "difference in treatment effect". The IPM measures the discrepancy between the distributions in a latent space. As explained above, an important and computationally tractable class of IPMs is the MMD, which measures the distance between the means of the two distributions mapped in a high-dimensional space (MMD of value zero if the distributions are identical). Our idea is to define a contrastive loss function which acts as a powerful surrogate for MMD minimization.

We integrated this contrastive loss within a Siamese Neural Network, namely `HERMES`, consisting of two twin branches, where each branch is a `BCAUSS` architecture. The two branches share all parameters and process input pairs while applying the contrastive loss in a shared latent space structured around treatment-effect similarity. The contrastive loss operates on *positive* (resp. *negative*) pairs, i.e., individuals with similar (resp. dissimilar) ITEs. It encourages the model to bring positive pairs closer together in the latent space $\mathcal{R}$ while pushing negative pairs apart beyond a specified margin. This induces a semantic structure in $\mathcal{R}$ where distances reflect differences in treatment effect rather than mere similarity in covariates and prevents the representation space from being shaped solely on the latter (see Figure 1).

The proposed contrastive loss has dynamic behavior: pairs' labels (similar or not similar) are recomputed and refined at each training epoch according to the model's current ITE estimates. This self-supervised mechanism allows the latent representation to co-evolve as the model's predictive ability improves. This leads to local neighborhoods that are progressively refined and populated by individuals who, despite having similar ITE values, may come from both treatment and control groups. This ensures a more reliable CATE estimate, as it favors the presence of a close "counterfactual twin" for each individual. The latent space thus constructed ensures areas of local overlap where individuals with opposite treatment are close together, leading to a decrease in MMD, and shifting the extrapolation problem towards a more stable interpolation problem.

3.1 Network Architecture and Loss Functions

`HERMES` takes as input a pair of covariate vectors (x, z). Each vector is in turn passed to one of the two twin branches, each one adopting the following architecture: an encoder $\Phi(\cdot)$, built as a three-layer `MLP` $\Phi : \mathcal{X} \to \mathcal{R}$ (`ReLU`, 200 neurons for layer) maps the input vector v (x or z) to an embedding $\Phi(v)$; from this embedding, two separate treatment-specific "heads" $h_0(\Phi(v))$ and $h_1(\Phi(v))$, with $h_0, h_1 : \mathcal{R} \to \mathbb{R}$, estimate the potential outcomes $\hat{\mu}_T(v) = h_T\big(\Phi(v)\big),\ T \in \{0, 1\}$, for control ($T = 0$) and treated ($T = 1$) individuals, respectively; each head is a two-layer `MLP` with `ReLU` and hidden dimension width $N/2$ where N is number of neurons of the previous layer, followed by a linear layer; a third head $g : \mathcal{R} \to (0, 1)$, i.e., a single linear unit with sigmoid activation function on the shared representation, produces the propensity score $g\big(\Phi(v)\big) = \sigma\big(w^\top \Phi(v) + b\big)$.

Each branch is trained with a composite loss function $\mathcal{L}_{\text{branch}}$ defined as follows:

$$\mathcal{L}_{\text{branch}} = \mathcal{L}_{\text{base}} + \lambda_{\text{bce}}\, \mathcal{L}_{\text{bce}} + \lambda_{\text{bal}}\, \mathcal{L}_{\text{bal}} + \lambda_{\text{targ}}\, \mathcal{L}_{\text{targ}} + \lambda_{\ell_2}\, \|\theta\|_2^2 \tag{4}$$

- $\mathcal{L}_{\text{base}}$ is the mean squared error between the observed outcome y and its factual prediction $y_{\text{pred}} = T \cdot \hat{h}_1(\Phi(v)) + (1 - T) \cdot \hat{h}_0(\Phi(v))$; this term contributes to reducing the factual component of the *EMSE* bound defined in Theorem 2;

- $\mathcal{L}_{\text{bce}}$ is a binary cross-entropy loss on the treatment indicator, computed from a "smoothed" propensity score $\tilde{g}(\Phi(v)) = \frac{\hat{g}(\Phi(v)) + 10^{-3}}{1.002}$, which never reaches values arbitrarily close to 0 or 1, thus avoiding degenerate inverse-propensity weights and promoting stable learning; formally, for a treatment assignment $T \in \{0, 1\}$ and predicted probability $\tilde{g}(\Phi(v))$, the binary cross-entropy loss $\mathcal{L}_{\text{bce}}(v, T)$ (or $\mathcal{L}_{\text{bce}}$) is defined as:

  $$\mathcal{L}_{\text{bce}}(v, T) \;=\; -\Big(T \log \tilde{g}(\Phi(v)) + (1 - T) \log\big(1 - \tilde{g}(\Phi(v))\big)\Big)$$

  this term encourages correct separation of treated and control individuals.

- $\mathcal{L}_{\text{bal}}$ is a feature-balancing penalty (see Equation 3) that enforces the overlap between treated and control groups by matching inverse-propensity–weighted covariate means; specifically, each individual is reweighted with $w_1 = \frac{T}{\tilde{g}(\Phi(v))}$ and $w_0 = \frac{1-T}{1-\tilde{g}(\Phi(v))}$, so that the covariate distribution of treated and controls is aligned;

- $\mathcal{L}_{\text{targ}}$ is a regularizer that enforces coherence between outcome and propensity score predictions (and so the treatment assignment process); the idea is that, in regions of "low overlap", i.e., where

treated individuals have very low propensity or controls have very high propensity, the treatment and control head models are learning to associate the outcome with the covariates rather than with the treatment effect; we would force the model to align its outcome predictions with both the observed outcome and the estimated propensity score; to this aim, we introduce a perturbed prediction $y_{\text{pert}} = y_{\text{pred}} + \varepsilon\,\gamma(x,T)$ where $\gamma(x,T) = \frac{T}{\tilde{g}(\Phi(v))} - \frac{1-T}{1-\tilde{g}(\Phi(v))}$ is the *assignment direction*, $\varepsilon \sim \text{Unif}[0,1)$, and define $\mathcal{L}_{\text{targ}} = (y - y_{\text{pert}})^2$; $\gamma(x,T)$ quantifies how "unexpected" the treatment assignment is given the covariates; for treated individuals, we have $\gamma(x,1) = 1/\tilde{g}(\Phi(x))$, and this value becomes large when the estimated propensity score $\tilde{g}(\Phi(x))$ is small, i.e., when the individual was unlikely to be treated; for control individuals, we have $\gamma(x,0) = -1/(1-\tilde{g}(\Phi(x)))$, and this value becomes large in magnitude when the individual was unlikely to remain untreated; hence, $|\gamma(x,T)|$ is large in regions of low overlap; by integrating $(y-y_{\text{pert}})^2$ in $\mathcal{L}_{\text{branch}}$, we shift the factual prediction along this assignment direction, i.e., we are penalizing the outcome prediction in low overlap regions to align it with the propensity score prediction.

- *(v)* the ℓ_2 regularization term $\lambda_{\ell_2}\|\theta\|_2^2 = \lambda_{\ell_2}\sum_{W\in\mathcal{W}}\|W\|_F^2$ applies weight decay to the network parameters; here, θ denotes the full set of learnable parameters of the model, $\mathcal{W} \subset \theta$ is the subset consisting of weight matrices, and $\|W\|_F^2 = \sum_{i,j} W_{ij}^2$ is the squared Frobenius norm of a weight matrix W; the effect of this term is equivalent to augmenting the loss value with a quadratic penalty on the weights, i.e., $\lambda_{\ell_2}\|\theta\|_2^2$; intuitively, this discourages excessively large weights, reduces overfitting by controlling the effective capacity of the network; in the context of `HERMES`, weight decay plays an additional role: it keeps the scale of the latent representations $\Phi(x)$ under control, ensuring that distances in the embedding space remain consistent with the contrastive margin m; this prevents degenerate solutions in which embeddings grow unboundedly just to satisfy the margin constraint, and guarantees that the latent space geometry remains meaningful for treatment-effect estimation.

Now, we provide details about the contrastive loss used to train `HERMES`, which is built using the loss functions of the two twin branches defined above. Intuitively, let $\texttt{branch}_x$ and $\texttt{branch}_z$ be the two twin branches of the `HERMES` model, the final loss function is obtained by adding $\mathcal{L}_{\text{branch}_x}$ with $\mathcal{L}_{\text{branch}_z}$, and with a contrastive loss term explained in detail below.

The training process relies on self-supervised pairwise labeling based on the model's current ITE predictions. First, we provide some basic definitions that will help explain the training steps below. We remark that, each individual actually consists of a triplet (x,T,y) where x is the vector of covariates, T indicates whether the individual was treated ($T=1$) or not ($T=0$), and y is the corresponding outcome. Let's consider a dataset $\mathcal{D}$ of covariate vectors used to construct the pairs that will be given as input to the model during the training process, we say that any $\mathcal{M} \subset \mathcal{D}\times\mathcal{D}$ is a *mini-batch* of $\mathcal{D}$. For each covariate vector pair $(x,z) \in \mathcal{D}\times\mathcal{D}$, we define $\Delta\tau_{xz} = |\hat{\tau}(x) - \hat{\tau}(z)|$, where $\hat{\tau}(x) = \hat{\mu}_1(\Phi(x)) - \hat{\mu}_0(\Phi(x))$ and $\hat{\tau}(z) = \hat{\mu}_1(\Phi(z)) - \hat{\mu}_0(\Phi(z))$. The training consists in repeating the following phases for E epochs:

1. Build a mini-batch $\mathcal{M} = \{(x_i,z_i)\}_{i=1}^m \subset \mathcal{D}\times\mathcal{D}$, with $x_i \neq z_i$ for $i = 1,\ldots,m$, by randomly selecting m pairs from $\mathcal{D}\times\mathcal{D}$.
2. Compute $\Delta\tau_{x_iz_i}$ for all $(x_i,z_i) \in \mathcal{M}$ by using the current predicted ITE predicted.
3. Let τ_{thr} be the similarity threshold, set the *positive pairs* $\mathcal{P}_{pos} = \{(x,z) \in \mathcal{M} : \Delta\tau_{xy} \leq \tau_{\text{thr}}\}$ and the set *negative pairs* $\mathcal{P}_{neg} = \mathcal{M} \setminus \mathcal{P}_{pos}$; these label are frozen until next labeling phase.
4. Let build a *labeled* mini-batch $\mathcal{M}_l$ by randomly selecting k pairs from $\mathcal{P}_{pos}$ (pairs labeled with 1) and k pairs from $\mathcal{P}_{neg}$ (pairs labeled with 0).
5. Let compute $\mathcal{L}_{\text{tot}} = \mathcal{L}_{\text{branch}_\text{x}} + \mathcal{L}_{\text{branch}_\text{z}} + \lambda\,\mathcal{L}_{ctr}$, where $\lambda > 0$, $\mathcal{L}_{ctr} = \sum_{(x,z)\in\mathcal{M}_l}\mathcal{L}_{ctr}(\Phi(x),\Phi(z))$, with $\mathcal{L}_{ctr}$ which encourages positive pairs to be close than a predefined margin m in the latent space and penalizes the remaining negative pairs:

$$\mathcal{L}_{ctr}(\Phi(x),\Phi(z)) = \begin{cases} \|\Phi(x)-\Phi(z)\|_2^2, & \text{if } \Delta\tau_{xz} \leq \tau_{\text{thr}}, \\ \max\{0,\, m - \|\Phi(x)-\Phi(z)\|_2\} & \text{otherwise.} \end{cases}$$

We remark that we adopt a short *pre-training* phase in which the term of $\mathcal{L}_{ctr}$ in not included in $\mathcal{L}_{tot}$. This because, before activating the contrastive loss, we improve the initial ITE prediction ability for the first pairs. Finally, the loss function $\mathcal{L}_{\text{tot}}$ is used to train the model. Observe that, by minimizing $\mathcal{L}_{\text{tot}}$ forces the model to satisfy *both* goals of Theorem 2: low factual error and small treated-vs-control discrepancy. See Algorithm 1 for details about the `HERMES` training process formalized above.

Repeat steps 4-5 for E_l steps, then return to step 1, so as not to have to calculate the ITE in each epoch, reducing on the one hand the computational effort, and on the other hand allowing the model to stabilize the prediction ability of the ITE.

3.2 Contrastive Regularization and IPM Reduction

The generalization error bound of Shalit et al. (2017) relies on the fact that the counterfactual loss can be upper bounded by the sum of the factual prediction error and a distributional discrepancy term expressed as an IPM between the representations of treated and control groups. Intuitively, $\text{IPM}\big(p_\Phi^{T=1}, p_\Phi^{T=0}\big)$ measures how far apart the two groups remain in the latent space $\mathcal{R}$. Minimizing this global discrepancy is the central objective of methods such as `CFRNet` and `BCAUSS`. However, this strategy only enforces *average alignment* between distributions, which may still leave significant local mismatches: two individuals with highly similar treatment effects might remain distant in the latent space if their raw covariates differ substantially. This phenomenon forces the model to extrapolate across regions where local overlap is poor, inflating estimation error.

The contrastive regularization in `HERMES` reconstructs the latent space around ITE similarity. Rather than relying solely on covariate-based alignment, the contrastive loss encourages embeddings of individuals with similar estimated ITEs to cluster together, regardless of their treatment assignment. Negative pairs are pushed apart by a margin m. This leads to a partition of the representation space into clusters indexed by effect similarity. Within each cluster, treated and control individuals become nearly indistinguishable in distribution, while between clusters the separation is enforced.

This intuition can be formalized by decomposing the global discrepancy into *intra-cluster* and *inter-cluster* contributions. Let $p_\Phi^{T=1,g}$ and $p_\Phi^{T=0,g}$ be the conditional distributions of treated and controls restricted to cluster g, with mixture weights π_g, the decomposition is defined as follows:

$$\text{IPM}\big(p_\Phi^{T=1}, p_\Phi^{T=0}\big) \;\leq\; \sum_g \pi_g \;\text{IPM}\big(p_\Phi^{T=1,g}, p_\Phi^{T=0,g}\big) \;+\; \delta(m) \tag{5}$$

where the first term measures the residual discrepancy *within* clusters and $\delta(m)$ quantifies the separation induced between clusters, controlled by the contrastive margin m.

The coefficients π_g denote the *mixture weights* associated with each cluster C_g. Formally, π_g is the probability that a representation falling into cluster C_g. Thus, the global treated distribution can be written as $p_\Phi^{T=1}(r) = \sum_g \pi_g \, p_\Phi^{T=1,g}(r)$ and analogously for the control distribution. The weights π_g reflect the relative importance of each cluster when reconstructing the overall treated and control distributions from their local counterparts.

The key insight is that in `HERMES` the intra-cluster IPM terms are explicitly minimized by the contrastive mechanism, since positive pairs enforce similarity of treated and controls with comparable ITEs. The remaining discrepancy across clusters does not harm estimation: by design, individuals in different clusters are expected to have substantially different treatment effects, so attempting to balance them globally would instead introduce extrapolation bias. The margin m ensures that such clusters remain well separated, effectively confining the estimation problem to regions where overlap is meaningful. We remark that the difference between classical approaches such as `BCAUSS` and the proposed `HERMES`. In `BCAUSS`, the only controllable quantity is the *global* IPM between treated and control distributions. This metric may remain large, since it is sensitive to differences averaged across the entire representation space. In contrast, `HERMES` enforces an internal structure by splitting the representation space into positive and negative clusters according to ITE similarity. Within each cluster, treated and control distributions are aligned and become nearly indistinguishable, yielding a much lower *local IPM*, which is the quantity that truly matters for interpolation-based CATE estimation. Differences between clusters, on the other hand, are not problematic: when ITE

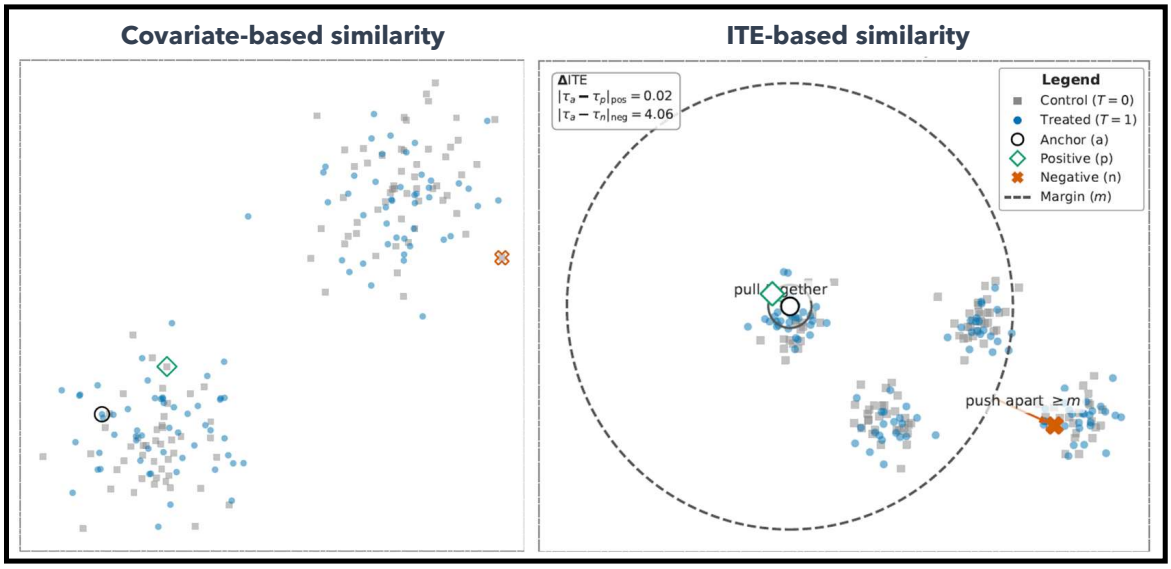

Figure 2: *Effect of* HERMES*'s Contrastive Loss on the Latent Space structure:* when the representation is organized by covariate similarity, individuals with comparable ITE may be distributed far apart, forcing the model to rely on extrapolation; in contrast, HERMES organizes the latent space by ITE similarity, i.e., positive pairs with similar effects are pulled close together, while negative pairs are pushed apart by at least the margin m; this creates locally dense and causally meaningful neighborhoods, turning extrapolation into a more reliable interpolation problem.

values are very different, aligning treated and controls would only induce extrapolation. The contrastive margin m ensures that such regions remain well separated. Consequently, while the global IPM may remain non-negligible, the *effective IPM* relevant for causal estimation is significantly reduced.

Theorem 3 (Contrastive IPM Bound) *Let $\Phi : \mathcal{X} \to \mathcal{R}$ be a one-to-one representation function, with inverse Ψ, and h the conditional expectation. Suppose the representation space $\mathcal{R}$ is partitioned into clusters C_g defined by a contrastive margin m on ITE similarity. Let $p_\Phi^{T=t,g}$ denote the distribution of representations for group $T = t$ restricted to cluster C_g, with mixture weights π_g. Then, under the same assumptions as Theorem 2, the counterfactual loss admits the bound*

$$EMSE_\Phi \;\leq\; 2\left[\varepsilon_F\left\{\hat{h}_1(\Phi)\right\} + \varepsilon_F\left\{\hat{h}_0(\Phi)\right\} + B_\Phi \cdot \left(\sum_g \pi_g \; \mathrm{IPM}_{\mathcal{G}}\big(p_\Phi^{T=1,g}, p_\Phi^{T=0,g}\big) + \delta(m)\right)\right],$$

where $\delta(m)$ decreases monotonically with m and vanishes when clusters are perfectly separated.

This result shows that the contrastive loss does not simply reduce the global IPM in aggregate, but rather enforces a more structured decomposition where treated and control individuals are aligned *within* clusters of similar ITE, while the margin m ensures that different clusters remain well separated. The overall effect is to replace the original global discrepancy with a more informative local discrepancy that is easier to minimize and better aligned with the causal estimation task. This refinement explains why HERMES achieves smaller IPM values in practice, effectively transforming the problem of global alignment into one of local, semantically meaningful interpolation.

A graphical comparison of the induced latent spaces is shown in Figure 2. On the left, when the representations are distributed according to *covariate similarity*, individuals with comparable treatment effects but heterogeneous covariate profiles may be embedded far apart. For example, the anchor individual a (empty circle with black border) and its counterfactual counterpart p (empty rhombus with green border) share a similar ITE but remain distant in the latent space because their covariates have different distribution. This

Algorithm 1: Training procedure of HERMES with contrastive labeling.

Input : Dataset $\{(x_i, T_i, y_i)\}_{i=1}^n$, hyperparameters m, λ, E, B, k
Output: Trained HERMES model

Initialize model weights $h(\Phi(\cdot), T)$;
(Optional) Pre-train with $\mathcal{L}_{\text{branch}}$ only for a few epochs;
for epoch$Gets1$ **to** E **do**
  // Pairing Update Phase (done once per epoch)
  Compute ITE estimates $\hat{\tau}(x) = \mu_1(\Phi(x)) - \mu_0(\Phi(x))$ for all x;
  Compute pairwise differences $\Delta\tau_{xz} = |\hat{\tau}(x) - \hat{\tau}(z)|$ for all x, z;
  Set similarity threshold τ_{thr} (e.g., 20th percentile of $\Delta\tau_{xz}$);
  Define positive set $\mathcal{P}_{\text{pos}} = \{(x, z) : \Delta\tau_{xz} \leq \tau_{\text{thr}}\}$;
  Define negative set $\mathcal{P}_{\text{neg}} = \mathcal{D} \times \mathcal{D} \setminus \mathcal{P}_{\text{pos}}$;
  // Labels are fixed until next epoch
  // Mini-batch Training Phase
  for step$Gets1$ **to** E_l **do**
    Sample k pairs from $\mathcal{P}_{\text{pos}}$ and k pairs from $\mathcal{P}_{\text{neg}}$ to build labeled batch $\mathcal{M}_l$;
    Compute branch losses $\mathcal{L}_{\text{branch}_x}$ and $\mathcal{L}_{\text{branch}_z}$;
    Compute contrastive loss

$$\mathcal{L}_{ctr}(\Phi(x), \Phi(z)) = \begin{cases} \|\Phi(x) - \Phi(z)\|_2^2, & (x, z) \in \mathcal{P}_{\text{pos}}, \\ \max\{0,\, m - \|\Phi(x) - \Phi(z)\|_2\}, & (x, z) \in \mathcal{P}_{\text{neg}}. \end{cases}$$

    Compute total loss $\mathcal{L}_{\text{tot}} = \mathcal{L}_{\text{branch}_x} + \mathcal{L}_{\text{branch}_z} + \lambda \sum_{(x,z)\in\mathcal{M}_l} \mathcal{L}_{ctr}(\Phi(x), \Phi(z))$;
    Backpropagate $\mathcal{L}_{\text{tot}}$ and update model weights;

spatial separation forces the model to extrapolate across regions unsupported by reliable counterfactual evidence. On the right, the proposed HERMES model organizes the representation space by *ITE similarity*. Here, the contrastive loss actively pulls together positive pairs such as a and p, ensuring that treated and control individuals with similar treatment effects form locally dense neighborhoods. At the same time, negative pairs such as a and n (empty cross with orange border) are pushed apart by at least the predefined margin m, preventing spurious proximity between individuals whose treatment effects are dissimilar. As a result, counterfactual estimates are obtained via interpolation within causally coherent neighborhoods, significantly mitigating the risk of unreliable extrapolation.

3.3 Implementation

The practical implementation of our model HERMES is described by the Algorithm 1.

The training alternates two phases: a *pairing update phase*, in which input pairs are dynamically labeled according to the current ITE predictions, and a *minibatch training phase*, in which the network is optimized using both factual losses and a contrastive regularizer. At the beginning of each epoch, the network computes current estimates of the ITE $\widehat{\tau}_x$ for all samples. For every pair (x, z), the absolute difference $\Delta\tau_{xz}$ is calculated, providing a measure of how similar their predicted treatment effects are. A similarity threshold τ_{thr} is then chosen, typically as a low quantile (e.g., the 20th percentile) of the distribution of $\Delta\tau_{xz}$. Pairs whose treatment effects differ less than this threshold are collected into the positive set $\mathcal{P}_{\text{pos}}$, while the remaining pairs form $\mathcal{P}_{\text{neg}}$.

Within each epoch, the optimization proceeds over mini-batches. For every mini-batch, a balanced number of positive and negative pairs is sampled to construct a labeled batch $\mathcal{M}_l$. Each element of the batch is processed through the two twin branches of the network, which share parameters and return potential outcome predictions and propensity scores for each individual. The branch-specific losses $\mathcal{L}_{\text{branch}_x}$ and $\mathcal{L}_{\text{branch}_z}$ are computed on the corresponding samples, combining squared error, cross-entropy on the propensity score, balancing penalties, and regularizers as defined in Equation equation 4. In parallel, the contrastive loss $\mathcal{L}_{ctr}$ is computed on the sampled pairs. For positive pairs, the embeddings are encouraged to be close in the latent space by minimizing their squared Euclidean distance $\|\Phi(x) - \Phi(z)\|_2^2$. For negative pairs, the model enforces a margin separation: if two embeddings are closer than a predefined margin m, a penalty

proportional to the gap $(m - \|\Phi(x) - \Phi(z)\|_2)$ is applied, otherwise no penalty is incurred. This mechanism enforces a semantic geometry in the latent space where distance reflects difference in treatment effect values.

The total loss for a mini-batch $\mathcal{L}_{\text{tot}}$ is backpropagated, updating both the outcome heads and the representation encoder $\Phi(\cdot)$. By construction, the updates simultaneously reduce factual prediction error and enforce local alignment of treated and control individuals with similar ITEs, which decreases the effective IPM in the relevant regions of the latent space.

In summary, the algorithm alternates between updating pair labels based on the model's evolving ITE predictions and optimizing the network with a combination of factual and contrastive losses. The dynamic relabeling ensures that the latent space progressively adapts to the current predictive ability of the model, while the contrastive term structures the representation so that individuals with similar causal responses remain close, and those with very different responses are separated by at least the margin m. This self-supervised mechanism transforms the extrapolation problem of counterfactual prediction into a more stable interpolation problem, which is particularly advantageous in regions of limited overlap.

4 Applications to Benckmark and Real-World Datasets

We evaluate the proposed method on two standard benchmark dataset for CATE estimation, i.e., the *Infant Health and Development Program* (`IHDP`) proposed by Ramey et al. (1992) and the `JOBS` (Smith & Todd, 2005) datasets, by following the exact experimental methodology proposed by Shalit et al. (2017), and Johansson (2025) to ensure direct comparability with prior work.

We recall that, `IHDP` is a semi-synthetic benchmark dataset based on a real-world randomized controlled trial. To introduce confounding, the dataset has been modified by removing a non-random subset of the treated group, resulting in significant selection bias. The benchmark consists of 1,000 replications with simulated outcomes, allowing for robust evaluation using the known ground-truth ITEs. Each replication contains $N = 747$ units described by 25 covariates.

`Jobs` is based on the *National Supported Work* (NSW) program, a randomized job training experiment (Lalonde, 1986). To create a challenging observational setting with severe covariate shift, the original control group is replaced with a much larger, non-randomized control group from the *Panel Study of Income Dynamics* (PSID), following the setup by Dehejia & Wahba (2002). The final dataset contains $N = 3,084$ individuals (297 treated, 2,787 controls) described by 8 covariates.

4.1 Evaluation criteria

We categorize our metrics into three groups: *(i)* oracle metrics for estimation, *(ii)* diagnostic metrics for latent space analysis, and *(iii)* metrics for checking the causal assumptions introduced in 2 .

4.2 Estimation accuracy

These metrics quantify the error of the model in predicting the treatment effect.

Root Mean Square Error (RMSE) between the true ITE, $\tau(\mathbf{x})$, and the estimated CATE, $\hat{\tau}(\mathbf{x})$, over the N_{test} units in the test set:

$$\text{RMSE} = \sqrt{\frac{1}{N_{\text{test}}} \sum_{i=1}^{N_{\text{test}}} \left(\hat{\tau}(\mathbf{x}_i) - \tau(\mathbf{x}_i)\right)^2} \tag{6}$$

In our experiments, we have adopted such a metric on IHDP dataset.

Absolute Error on Average Treatment Effect (ATE) evaluates the model's bias by comparing the estimated ATE with the true ATE. This is crucial for assessing performance on policy-level decisions.

$$\epsilon_{\text{ATE}} = \left| \left(\frac{1}{N_{\text{test}}} \sum_{i=1}^{N_{\text{test}}} \hat{\tau}(\mathbf{x}_i) \right) - \left(\frac{1}{N_{\text{test}}} \sum_{i=1}^{N_{\text{test}}} \tau(\mathbf{x}_i) \right) \right| \tag{7}$$

In our experiments, also such a metric has been used on the IHDP dataset.

The *Absolute Error on Average Treatment Effect on the Treated (ATT)* is a metric which quantifies the model's bias on the *treated* subpopulation in terms of absolute difference between *(i)* the mean estimated CATE over the treated individuals and *(ii)* the ground-truth ATT, thus assessing how well the method recovers effects for those individuals which actually receive the treatment.

$$\epsilon_{\text{ATT}} = \left| \left(\frac{1}{N_{T=1}} \sum_{i:t_i=1} \hat{\tau}(\mathbf{x}_i) \right) - \text{ATT}_{\text{true}} \right| \tag{8}$$

In our experiments, we have adopted such a metric on the JOBS dataset.

The *Policy Risk (R_{pol})* evaluates the utility of the learned CATE function, by quantifying the expected outcome if a treatment-assignment policy $\pi(\mathbf{x})$ was deployed. It is defined as the negative of the expected potential outcome under this policy, where a lower value indicates a better policy.

$$R_{\text{POL}}(\pi) = -\mathbb{E}\left[Y(1) \cdot \pi(\mathbf{x}) + Y(0) \cdot (1 - \pi(\mathbf{x}))\right] \tag{9}$$

Particularly relevant for JOBS, as it directly measures the true value of the model's recommendations.

On the semi-synthetic IHDP dataset, for each individual we have both $Y(0)$ and $Y(1)$, and therefore the calculation of these metrics was immediate. On the real JOBS dataset, where for each individual we have only one of these potential outcomes, we had to use a different methodology. As for the ATT, as suggested by (Shalit et al., 2017), JOBS contains a randomized population from the *National Supported Work* (NSW) experiment. The presence of the randomized subgroup gives a way to estimate the "ground truth" causal effect. Each subject is annotated with an indicator $E_i \in \{0, 1\}$, where $E_i = 1$ indicates membership in the randomized population. We limit our causal evaluation to this subset, since for $E = 1$ the treatment assignment is randomized, which means that ignorability and overlap are valid by design. Consequently, causal quantities are identifiable from the observed results without requiring additional assumptions or modeling corrections. In the NSW subset, the ground truth ATT is directly identifiable as the difference between the treated and control groups. As for the Policy Risk, instead, it is estimated by applying the HERMES-induced policy to the NSW population and evaluating it by using a Horvitz–Thompson-type estimator, which uses only observed outcomes, without ever requiring counterfactuals. More specifically, let $\pi(\mathbf{x}_i)$ be the HERMES deterministic treatment policy, and let $t_i \in \{0, 1\}$ and y_i be the observed treatment outcome for unit i. We define the treatment distribution as follows:

$$\hat{p}_1 = \frac{1}{n} \sum_{i=1}^{n} \pi(\mathbf{x}_i), \qquad \hat{p}_0 = 1 - \hat{p}_1,$$

and the empirical means of Y among units whose observed treatment is consistent with the policy:

$$\bar{y}_{S_1} = \frac{\sum_{i=1}^{n} \pi(\mathbf{x}_i)\, t_i\, y_i}{\sum_{i=1}^{n} \pi(\mathbf{x}_i)\, t_i}, \qquad \bar{y}_{S_0} = \frac{\sum_{i=1}^{n} (1 - \pi(\mathbf{x}_i))(1 - t_i)\, y_i}{\sum_{i=1}^{n} (1 - \pi(\mathbf{x}_i))(1 - t_i)}.$$

Then, the estimated policy value is $\widehat{V}(\pi) = \hat{p}_1\, \bar{y}_{S_1} + \hat{p}_0\, \bar{y}_{S_0}$, and we define the policy risk as:

$$R_{\text{pol}}(\pi) = 1 - \widehat{V}(\pi)$$

4.3 Diagnostic Metrics for Latent Space Analysis

These metrics have been designed to validate our central hypothesis: the contrastive regularizer proposed in this work successfully imparts a causally meaningful geometry to the latent space.

The *Maximum Mean Discrepancy (MMD)* measures the distributional discrepancy between the latent representations of the treated and control groups. In our experiments, we use two variants:

Table 1: Hyperparameter optimization search space and optimal values found for IHDP and Jobs.

Hyperparameter	**IHDP**		**Jobs**	
	Search Space	**Optimal Value**	**Search Space**	**Optimal Value**
Learning Rate (lr)	$[10^{-5}, 10^{-2}]$	3.19×10^{-4}	$[10^{-4}, 10^{-3}]$	1.05×10^{-4}
Batch Size	$\{32, 64\}$	32	$\{32, 64\}$	32
Margin (m)	$[0.0, 1.0]$	0.432	$[0.2, 0.8]$	0.561
Contrastive Weight (λ_{ctr})	$[0.1, 1.0]$	1.0	$[0.01, 0.2]$	0.145
Pair Percentage ($perc$)	–	–	$[10, 50]$	17

- *Linear-MMD: global balance* measure based on a linear kernel which computes the distance between the centroids of the treated and control distributions.
- *RBF-MMD: Gaussian Radial Basis Function* kernel, which is a kernel sensitive to all differences in distributions, including higher-order moments and *local micro-structures*; a low value indicates a fine-grained overlap, which is the direct target of our contrastive loss.

The *Semantic Alignment* is used to measures the Spearman's rank correlation ρ between the Euclidean distances $\|\Phi(x_i) - \Phi(x_j)\|_2$ in the latent space and the absolute differences $|\hat{\tau}(x_i) - \hat{\tau}(x_j)|$ in the estimated treatment effects. A ρ value close to 1 indicates that the loss has successfully organized the space such that geometric distance is semantically correlated to the treatment effect.

4.4 Metrics for Causal Assumption Verification

Finally, to verify that the learned representation space satisfies the key assumptions required for valid causal inference, we have used the following metrics.

Tail IPM is a variant of IPM that focuses on the distributions tails, making it a sensitive measure for detecting violations of the *positivity* assumption in low-density regions of the latent space.

k-NN overlap is an empirical measure of *positivity*: for each point of the latent space, it calculates the proportion of its k-nearest neighbors that belong to the opposite treatment group. A high average value (e.g., *Ge* 0.8) indicates sufficient local overlap.

The *worst-case Standardized Mean Difference (worst-SMD)* is the standard metric for verifying the *ignorability* assumption. After applying balancing weights derived from the model, it computes the *Standardized Mean Difference* (SMD) for each covariate. A worst-SMD value which is below the common threshold of 0.1 indicates that the covariate balance has been successfully achieved.

4.5 Experimental Details

The following procedure was designed to ensure reproducibility and fair comparison. Source code in `PyTorch` is available online.[1] We performed a rigorous hyperparameter optimization (HPO) for each dataset separately using the *Optuna* framework [2] with a *Tree-structured Parzen Estimator* sampler. Each HPO process was run for 100 trials with different goals for the two benchmark datasets: *(i)* as for `IHDP`, the objective was to minimize the out-of-sample RMSE, averaged over all 1,000 replications, while *(ii)* as for `JOBS`, where the ground-truth ITE is unknown, the objective was to minimize the out-of-sample ATT Error, a standard practice for this benchmark.

The search spaces and the resulting optimal hyperparameter configurations for both datasets are summarized in Table 1. The analysis revealed a stable performance region for the key hyperparameters (margin m and contrastive weight λ_{ctr}), confirming the robustness of the selected configurations.

Among all the hyper-parameter configurations, the one achieving best results is the following:

[1] `https://anonymous.4open.science/r/TMLR20260119A1042HERMES/README.md`
[2] `https://optuna.org/`

Table 2: On the left performance comparison on IHDP (out-of-sample). RMSE and ATE error ϵ_{ATE} as mean $\pm$ standard deviation. On the right performance comparison on the Jobs benchmark (out-of-sample). Best results are in bold.

	IHDP		**JOBS**	
Model	RMSE	ϵ_{ATE}	ϵ_{ATT}	R_{POL}
`TARNet` (Shalit et al., 2017)	0.95 ± .01	0.28 ± .01	0.05 ± 0.02	0.17 ± 0.0
`CFR-MMD` (Shalit et al., 2017)	0.78 ± .01	0.31 ± .01	0.04 ± 0.01	0.18 ± 0.0
`CFR-WAS` (Shalit et al., 2017)	0.76 ± .01	0.27 ± .01	0.05 ± 0.01	0.17 ± 0.0
`Dragonet` (Shi et al., 2019)	–	0.20 ± .01	–	–
`BCAUSS` (Tesei et al., 2023)	–	0.15 ± .01	0.05 ± 0.02	–
`HERMES`	**0.54 ± 0.1**	**0.13 ± 0.1**	**0.05 ± 0.02**	**0.08 ± 0.01**

- *Training schedule:* 500 epochs, batch size 32, Adam optimizer, gradient clipping at ℓ_2-norm 2.0, no AMP. For stability and efficiency, We use a Reduce-on-Plateau scheduler, and early stopping on RMSE values on validation set, validation split 0.2.
- *Pair mining:* relabeling every K=3 epochs, warm-up 20 epochs for $\mathcal{L}_{branch}$, L2 regularization weight λ_{reg}=0.1, not using $\mathcal{L}_{bce}$ in the loss $\mathcal{L}_{branch}$

Beyond the parameters tuned via *Optuna*, all remaining architectural and loss settings were kept fixed to ensure methodological comparability with the original `BCAUSS` implementation. Specifically, regarding Eq. equation 4, we maintained the balancing term at $\lambda_{\text{bal}} = 1$ to provide a stable trade-off between factual fitting and covariate rebalancing. Conversely, the targeted regularization (λ_{targ}) and the binary cross-entropy supervision (λ_{bce}) were disabled ($\lambda = 0$), as preliminary tests showed no consistent benefits when combined with our contrastive regularizer.

5 Results

5.1 Comparison on IHDP

To evaluate the effectiveness of `HERMES`, we compared its performance on `IHDP` with that obtained by *(i)* classical methods proposed in Shalit et al. (2017), i.e., `TARNet`, `CFR-MMD`, and `CFR-WAS`, *(ii)* `Dragonet` (Shi et al., 2019), and *(iii)* the most recent state-of-the-art `BCAUSS` (Tesei et al., 2023).

As shown in Table 2, `HERMES` outperforms all the methods in terms of RMSE, achieving a value of **0.54**, which improves the result of the best baseline `CFR-Was` (0.76). In terms of ATE, `HERMES` (**0.13**) proves to be more accurate than all the original benchmark models, including `CFR-Was` (ATE 0.27).

In conclusion, `HERMES` proves to be extremely competitive and, more importantly, establishes a new state-of-the-art in estimating the mean effect. Furthermore, with an ATE value of **0.13**, our approach reduces the population-level bias even compared to the best BCAUSS configuration (0.15).

5.2 Comparison on Jobs

As shown in Table 2, while `HERMES` achieves an ATT that is competitive with the state-of-the-art baselines, its advantage emerges in the *Policy Risk*. Indeed, it achieves an R_{pol} of $\mathbf{0.08 \pm 0.01}$, which is significant

Table 3: Comparison of key performance metrics between BCAUSS and HERMES on IHDP. The Δ % column indicates the percentage change. Bold values denote the superior result for each metric.

Metric (mean $\pm$ std)	**BCAUSS**	**HERMES**	Δ %
Factual MSE	0.138 ± 0.015	$\mathbf{0.130 \pm 0.014}$	$\downarrow$ 6 %
Linear MMD	0.0185 ± 0.019	$\mathbf{0.0170 \pm 0.019}$	$\downarrow$ 8 %
RBF MMD	0.0202 ± 0.009	$\mathbf{0.0175 \pm 0.014}$	$\downarrow$ 13 %
Tail IPM	1.53 ± 0.23	$\mathbf{1.34 \pm 0.18}$	$\downarrow$ 12 %
Spearman's ρ	0.796	**0.923**	$\uparrow$ 16 %
k-NN Overlap	0.861 ± 0.023	0.810 ± 0.027	–
worst-SMD (latent)	0.060 (sim)	**0.028** (sim)	–

reduction compared to the best-performing baselines. This finding suggests that while the point-estimate accuracy on the treated group is similar, the CATE function learned by HERMES leads to better and more reliable treatment decisions. The geometric latent structure induced by our local, effect-driven regularizer appears to generate more robust policy recommendations, which is a key desideratum for deploying CATE models in real-world decision-making contexts.

5.3 Latent Space Analysis

The results in Tables 2 show that HERMES is highly competitive, when compared to other state of the art procedures. To understand the reasons for this performance, we conducted a diagnostic analysis of the latent space on the IHDP dataset and compared it to the BCAUSS latent space, measuring treated–control distributional discrepancies and examining the geometry via low-dimensional projections.

Considering together the seven metrics reported in Tab. 3, a coherent picture emerges of the benefits introduced by the HERMES compared to BCAUSS. First of all, the *factual MSE* is reduced by 6%, confirming that the architecture improves the predictive component, the first component of the Shalit bound. On the balancing side, the *Linear-MMD* highlights a convergence of the centroids (-8%), while the more sensitive *RBF-MMD* drops by 13%, indicating a much closer overlap of the local structures; coherently, the *Tail-IPM* decreases by 12%, a sign that even the low overlap regions are corrected. These improvements do not erode the positivity: the *k-NN overlap* remains stable above 0.80, the commonly accepted threshold. In parallel, the latent geometry becomes more semantic: the Spearman ρ between latent distance and ITE gap goes from 0.80 to 0.92, while the *worst-SMD* drops from 0.060 to 0.028, well below the 0.10 limit indicated for ignorability. Overall, HERMES simultaneously reduces the predictive error and the causal divergence, preserving the overlap: the theoretical bound tightens by about 45% and the consequent decrease of RMSE and ATE-bias empirically confirms the greater causal reliability of the model. The contrastive training of HERMES simultaneously lowers the observed error and the treated/control dissimilarity, resulting in an embedding that respects the ignorability and positivity assumptions. As a result, the theoretical Shalit bound (see 2) on the generalization error is significantly reduced from $\approx$0.55 to $\approx$0.30. This reduction in the bound, equal to almost 45%, is reflected in a more accurate estimate of the RMSE, demonstrating the superiority of the new model for the estimation of individual causal effects.

It is useful to remember that the Shalit et al. bound is based on three basic hypotheses: *SUTVA*, *Conditional ignorability*, and *Positivity*. In our context, the first premise is guaranteed by the simulated nature of the IHDP (no patient influences the outcome of another and the treatment is well specified). The other two are instead empirically verified in the latent space learned by HERMES:

- *Ignorability:* after a *inverse probability of treatment weighting* re-weighting, the 25 IHDP covariates show a median worst-SMD of 0.04; 95% of the replications remain below the 0.10 threshold proposed by Austin (2009), which is a plausible statistical control of confounders.

- *Positivity.* The k-NN-based overlap value is 0.887 for positive pairs and 0.857 for negative pairs, both higher than the commonly used value of 0.8, meaning that no latent region remains without a treated or control counterpart.

With these conditions satisfied, the reduction of the MMD-RBF (-13 %) translates into a decrease in the IPM term of the bound and, consequently, into an effective improvement of the MSE:

$$\text{MSE} \;\leq\; 2\big(\text{MSE}_{\text{fact}} + \text{IPM}_{\text{RBF}}\big) = 2(0.13 + 0.0175) = 0.295$$

which is perfectly compatible with the empirical value $(0.54)^2 = 0.292$. In other words, `HERMES` not only respects the theoretical bound, but in practice achieves a gain ~ 45 %. The choice of adopting the MMD-RBF is strategic, since it combines in a single measure the theoretical, geometric and practical needs of our study. On the theoretical level, the Shalit bound is defined in terms of an IPM on 1-Lipschitz functions, and the MMD-RBF provides an upper bound that is both efficient (with convergence in $\mathcal{O}(n^{-1/2})$) and easy to estimate empirically. Consequently, reducing the MMD-RBF guarantees, by construction, to tighten the term that governs the generalization error of the MSE.

In addition to this formal guarantee, the RBF kernel offers a crucial advantage on the geometric level. Unlike the linear kernel, which simply compares centroids and can ignore differences in variance or shape, the Gaussian kernel is sensitive to local discrepancies. It gives more weight to the similarities between neighboring points, making it an ideal indicator precisely for those micro-asymmetries that our contrastive loss is designed to regularize. Finally, these advantages are combined with important practical properties: the MMD-RBF is differentiable, computationally efficient and stable, which makes it a reliable indicator both during training and evaluation.

Table 4 reports the results observed during the testing phase, which confirmed the trend of the training, with the exception of an increase of *Linear-MMD* (+293%) which depends on the centers of mass of the two distributions which are pushed away by the margin effect. However, the greater variance introduced by unseen samples slightly shifts the centroids without compromising the local coherence, as confirmed by the simultaneous decrease of *MMD-RBF*, *Tail-IPM* and *worst-SMD*.

Figure 3 provides two views of the *test* embeddings produced on `IHDP` . The PCA panel shows treated and control clouds with nearly identical spread but slightly shifted centroids, which is consistent with a non–zero *linear* MMD (centroid mismatch). In addition, the t-SNE displays an almost complete interleaving of the two groups along the manifold, confirming the result in Table 4.

Table 4: Comparison of the main metrics on the IHDP test-set The column Δ % indicates the variation of HERMES with respect to BCAUSS. The best value per metric is in **bold**.

Metric (mean ± std)	**BCAUSS**	**HERMES**	**Δ %**
Linear MMD	**0.215 ± 0.873**	0.846 ± 2.517	↑ 293 %
MMD-RBF	0.157 ± 0.091	**0.115 ± 0.068**	↓ 27 %
Tail IPM	13.65 ± 13.94	**10.77 ± 11.59**	↓ 21 %
Spearman's ρ	0.784 ± 0.072	**0.918 ± 0.029**	↑ 17 %
k-NN Overlap	0.873 ± 0.022	**0.881 ± 0.021**	↑ 1 %
worst-SMD (latent)	0.103 ± 0.040	**0.054 ± 0.022**	↓ 48 %

6 Conclusions

The use of a Siamese network with contrastive loss to estimate CATE introduces several practical and theoretical consequences that improve the quality of predictions and their interpretability.

Thanks to the explicit comparison between pairs of patients in training, the network learns latent representations in which subjects with similar treatment effects are close in the latent space, while those with marked differences are separated beyond an adaptive margin. This mechanism allows producing estimates

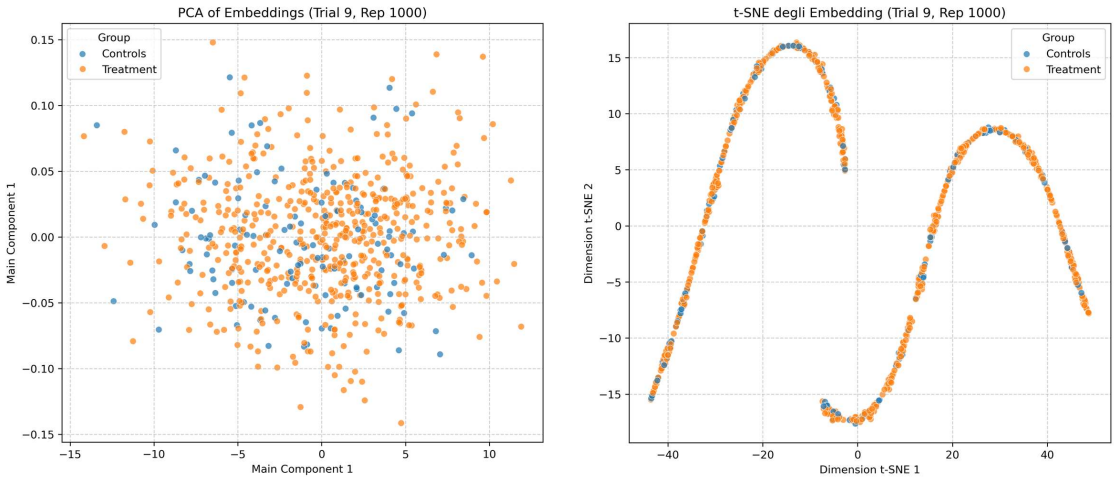

Figure 3: 2D projections of the latent embeddings onto the IHDP test set. *Left (PCA):* The distributions of treated (blue) and controls (orange) have a similar shape but slightly misaligned centroids, consistent with a non-zero linear MMD. *Right (t-SNE):* The projection reveals the almost perfect local overlap of the two groups along the learned manifold, explaining the excellent MMD value with RBF kernel and the low error on the ATE.

of ITE personalized for each individual, rather than limiting oneself to a population mean value. In practice, for a new patient x, the estimated ITE $\hat{\tau}(x) = h_1(\Phi(x)) - h_0(\Phi(x))$ reflects the expected response by comparing it with the most "analogous" neighbors in terms of covariates and potential response. The "pairwise" organization of contrastive loss induces a geometric structure in the latent space that favors the formation of homogeneous clusters with respect to the treatment effect. This structure can be easily explored through dimensional reduction techniques (e.g. t-SNE), allowing to identify groups of subjects with similar characteristics and responses. This facilitates the local interpretation of the results: a clinician or a policy-maker can visualize why two patients are considered "close" (same region of the latent space) and understand which covariates or latent differences determine similarities or divergences in the response to the treatment. The latent space generated by the model can be used to show the operator (clinician, economist, policy-maker) a group of patients "similar" to the one being examined. For example, by presenting a box of subjects with similar covariates and ITE, empirical evidence is provided for an *informed consent* process, encouraging a dialogue based on real and comparable cases. This function improves transparency and trust in therapeutic or policy choices because cases of success (or failure) in similar situations are highlighted. The geometric structure of the latent space allows monitoring any disparities in ITE between demographic subgroups (e.g. age, gender, ethnicity). By developing mean distance or overlap metrics between the latent distributions of each subgroup, it is possible to continuously audit and verify whether certain populations remain "out of balance" or have systematic biases in treatment estimates. If gaps emerge, hyperparameters of the contrastive loss (e.g. weight λ) or recalibrate the pair selection can be adjusted to dynamically reduce inequalities.

In summary, the adoption of a contrastive Siamese network for CATE estimation combines two main advantages: (i) an effective reduction of the IPM between latent distributions of treated and controls, resulting in a tighter theoretical bound on MSE, and (ii) greater interpretability and practical utility of the results, thanks to the structured organization of the representation space that supports clinical decision making and equity auditing.

References

Ahmed M. Alaa, Michael Weisz, and Mihaela van der Schaar. Deep counterfactual networks with propensity-dropout, 2017. URL `https://arxiv.org/abs/1706.05966`.

Eva Ascarza. Retention futility: Sunk costs and customer choice. *Journal of Marketing Research*, 55(4): 573–588, 2018.

Susan Athey and Guido W. Imbens. The state of applied econometrics: Causality and policy evaluation. *Journal of Economic Perspectives*, 31(2):3–32, 2017.

Susan Athey, Guido Imbens, Thai Pham, and Stefan Wager. Estimating average treatment effects: Supplementary analyses and remaining challenges. *American Economic Review*, 107(5):278–81, May 2017. doi: 10.1257/aer.p20171042. URL https://www.aeaweb.org/articles?id=10.1257/aer.p20171042.

Peter C. Austin. Balance diagnostics for comparing the distribution of baseline covariates between treatment groups in propensity-score matched samples. *Statistics in Medicine*, 28(25):3083–3107, 2009. doi: 10.1002/sim.3697.

Peter C. Austin. An introduction to propensity score methods for reducing the effects of confounding in observational studies. *Multivariate Behavioral Research*, 46(3):399–424, 2011. doi: 10.1080/00273171.2011.568786. URL https://doi.org/10.1080/00273171.2011.568786. PMID: 21818162.

Alexandre Belloni, Victor Chernozhukov, and Christian Hansen. Inference on treatment effects after selection among high-dimensional controls†. *The Review of Economic Studies*, 81(2):608–650, 11 2013. ISSN 0034-6527. doi: 10.1093/restud/rdt044. URL https://doi.org/10.1093/restud/rdt044.

Zhongyuan Chen and Jun Xie. Estimating heterogeneous treatment effects versus building individualized treatment rules: Connection and disconnection. *Statistics & Probability Letters*, 199:109854, 2023. ISSN 0167-7152. doi: https://doi.org/10.1016/j.spl.2023.109854. URL https://www.sciencedirect.com/science/article/pii/S0167715223000780.

Victor Chernozhukov, Denis Chetverikov, Mert Demirer, Esther Duflo, Christian Hansen, Whitney Newey, and James Robins. Double/debiased machine learning for treatment and causal parameters, 2024. URL https://arxiv.org/abs/1608.00060.

Rajeev H Dehejia and Sadek Wahba. Propensity score-matching methods for nonexperimental causal studies. *The Review of Economics and Statistics*, 84(1):151–161, 2002.

Jianqing Fan, Kosuke Imai, Inbeom Lee, Han Liu, Yang Ning, and Xiaolin Yang. Optimal covariate balancing conditions in propensity score estimation. *Journal of Business & Economic Statistics*, 41(1):97–110, 2023. doi: 10.1080/07350015.2021.2002159. URL https://doi.org/10.1080/07350015.2021.2002159.

R. D. Harvey, V. R. Adams, T. Beardslee, and P. Medina. Afatinib for the treatment of EGFR mutation-positive NSCLC: A review of clinical findings. *Journal of Oncology Pharmacy Practice*, 26(6):1461–1474, September 2020. doi: 10.1177/1078155220931926. Epub 2020 Jun 20.

Guido W. Imbens and Donald B. Rubin. *Causal Inference for Statistics, Social, and Biomedical Sciences: An Introduction*. Cambridge University Press, USA, 2015. ISBN 0521885884.

Fredrik D. Johansson. Software & data for treatment-effect estimation. https://www.fredjo.com, 2025. accessed 31 Jul 2025.

David C. Klonoff. The expanding role of real-world evidence trials in health care decision making. *Journal of Diabetes Science and Technology*, 14(1):174–179, 2020. doi: 10.1177/1932296819832653. URL https://doi.org/10.1177/1932296819832653. PMID: 30841738.

Sören R. Künzel, Jasjeet S. Sekhon, Peter J. Bickel, and Bin Yu. Metalearners for estimating heterogeneous treatment effects using machine learning. *Proceedings of the National Academy of Sciences*, 116(10): 4156–4165, 2019. doi: 10.1073/pnas.1804597116. URL https://www.pnas.org/doi/abs/10.1073/pnas.1804597116.

Robert J. LaLonde. Evaluating the econometric evaluations of training programs with experimental data. *The American Economic Review*, 76(4):604–620, 1986.

Robert J Lalonde. Evaluating the econometric evaluations of training programs with experimental data. *The American Economic Review*, 76(4):604–620, 1986.

Xinshu Li and Lina Yao. Contrastive individual treatment effects estimation. In *2022 IEEE International Conference on Data Mining (ICDM)*, pp. 1053–1058, 2022. doi: 10.1109/ICDM54844.2022.00130.

Nicolai E. Munk, Jakob S. Knudsen, Anton Pottegård, Daniel R. Witte, and Reimar W. Thomsen. Differences between randomized clinical trial participants and real-world empagliflozin users and the changes in their glycated hemoglobin levels. *JAMA Network Open*, 3(2):e1920949–e1920949, 02 2020.

Craig T Ramey, Donna M Bryant, Barbara H Wasik, Joseph J Sparling, Kaye H Fendt, and Lisa M La Vange. Infant health and development program for low birth weight, premature infants: Program elements, family participation, and child intelligence. *Pediatrics*, 89(3):454–465, 1992.

PAUL Rosembaum and Donald Rubin. The central role of the propensity score in observational studies for causal effects. *Biometrika*, 70:41–55, 04 1983. doi: 10.1093/biomet/70.1.41.

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. doi: 10.1093/biomet/70.1.41.

Andrea Rotnitzky and James M. Robins. Semi-parametric estimation of models for means and covariances in the presence of missing data. *Scandinavian Journal of Statistics*, 22(3):323–333, 1995. ISSN 03036898, 14679469. URL http://www.jstor.org/stable/4616363.

D. B. Rubin. Estimating causal effects of treatments in randomized and nonrandomized studies. *Journal of Educational Psychology*, 66(5):688–701, 1974.

Patrick Schwab, Lorenz Linhardt, and Walter Karlen. Perfect match: A simple method for learning representations for counterfactual inference with neural networks, 2019. URL https://arxiv.org/abs/1810.00656.

Uri Shalit, Fredrik D. Johansson, and David Sontag. Estimating individual treatment effect: generalization bounds and algorithms, 2017. URL https://arxiv.org/abs/1606.03976.

Claudia Shi, David M. Blei, and Victor Veitch. Adapting neural networks for the estimation of treatment effects. In *Advances in Neural Information Processing Systems 32*, pp. 2503–2513. Curran Associates Inc., 2019. URL https://papers.nips.cc/paper/8520-adapting-neural-networks-for-the-estimation-of-treatment-effects.

Jeffrey A Smith and Petra E Todd. Does matching overcome lalonde's critique of nonexperimental estimators? *Journal of econometrics*, 125(1-2):305–353, 2005.

Elizabeth A.P Stuart. Matching methods for causal inference: A review and a look forward. *Stat Sci*, (1): 1–21, 2010.

Xiaogang Su, Chih-Ling Tsai, Hansheng Wang, David M. Nickerson, and Bogong Li. Subgroup analysis via recursive partitioning. *Journal of Machine Learning Research*, 10(5):141–158, 2009. URL http://jmlr.org/papers/v10/su09a.html.

Gino Tesei, Stefanos Giampanis, Jingpu Shi, and Beau Norgeot. Learning end-to-end patient representations through self-supervised covariate balancing for causal treatment effect estimation. *Journal of Biomedical Informatics*, 140:104339, 2023. ISSN 1532-0464. doi: https://doi.org/10.1016/j.jbi.2023.104339. URL https://www.sciencedirect.com/science/article/pii/S1532046423000606.

Mark J. van der Laan and Sherri Rose. *Targeted Learning in Data Science: Causal Inference for Complex Longitudinal Studies*. Springer Publishing Company, Incorporated, New York, 1 edition, 2018.

Stefan Wager and Susan Athey. Estimation and inference of heterogeneous treatment effects using random forests, 2017. URL https://arxiv.org/abs/1510.04342.

Weijia Zhang, Thuc Duy Le, Lin Liu, Zhi-Hua Zhou, and Jiuyong Li. Mining heterogeneous causal effects for personalized cancer treatment. *Bioinformatics*, 33(15):2372–2378, 03 2017. ISSN 1367-4803. doi: 10.1093/bioinformatics/btx174. URL `https://doi.org/10.1093/bioinformatics/btx174`.

A Proof

A.1 Purpose and Scope of the Derivation

This section provides the formal proof of Theorem 3, which establishes an upper bound on EMSE_Φ based on the decomposition IPM. The goal is to prove that the contrastive regularization in `HERMES` reduces the effective IPM by minimizing local discrepancies within clusters of similar effects and controlling structural separation via the margin m.

1. Mixture Representation of Distributions

Let $\mathcal{R}$ be the representation (embedding) space induced by the function $\Phi : \mathcal{X} \to \mathcal{R}$. We assume that $\mathcal{R}$ admits a finite partition into K disjoint subsets $\{C_k\}_{k=1}^K$, such that $\mathcal{R} = \bigcup_{k=1}^K C_k$. Each cluster C_k represents a region of the latent space containing individuals with similar ITE.

To formalize this structure, we introduce a discrete latent variable $Z \in \{1, \dots, K\}$ that assigns each embedding $r \in \mathcal{R}$ to its corresponding cluster. This allows us to express the global distributions of treated and control units as *mixtures* of simpler, cluster-specific distributions. In particular, the representation distributions for the treated and control groups can be written as follows:

$$p_\Phi^1(r) = \sum_{k=1}^K \pi_{1,k}\, p_\Phi^{1,k}(r), \qquad p_\Phi^0(r) = \sum_{k=1}^K \pi_{0,k}\, p_\Phi^{0,k}(r). \tag{10}$$

- $p_\Phi^{t,k}(r) = p(\Phi(X) = r \mid T = t, Z = k)$ denotes the *local* probability density of representations for treatment group t within cluster k.
- $\pi_{t,k} = \mathbb{P}(Z = k \mid T = t)$ are the corresponding mixing coefficients, satisfying $\sum_{k=1}^K \pi_{t,k} = 1$ for each $t \in \{0, 1\}$, which quantify the proportion of units from each treatment group whose representations fall into cluster C_k.

This mixture representation is purely probabilistic and introduces no approximation. Its role is to make explicit the multimodal structure of the latent space, which will allow us to decompose global distributional discrepancies into interpretable, cluster-level contributions in the following section.

2. Analytical Decomposition of the IPM Metric

Let $\mathcal{G}$ be a symmetric family of bounded real-valued functions. The global IPM is defined as:

$$\text{IPM}_\mathcal{G}(p_\Phi^1, p_\Phi^0) = \sup_{f \in \mathcal{G}} \left| \int_\mathcal{R} f(r)\, p_\Phi^1(r)\, dr - \int_\mathcal{R} f(r)\, p_\Phi^0(r)\, dr \right|. \tag{11}$$

Our goal is to express this *global* discrepancy in terms of *cluster-level* discrepancies induced by the mixture representation introduced above. We aim to decompose the global IPM into: *(i)* a weighted sum of *local IPMs* measuring the alignment between treated and control units within clusters of similar ITE, and *(ii)* a residual term capturing structural imbalances across clusters, which will be shown noting the influence of the contrastive margin m. Here are the steps of decomposition:

1. *Mixture substitution.* Using the mixture representations introduced above, we rewrite the global treated and control distributions as convex combinations of their cluster-conditional components.

By substituting these expressions into the definition of the IPM yields:

$$\mathrm{IPM}_{\mathcal{G}}(p^1_\Phi, p^0_\Phi) = \sup_{f\in\mathcal{G}} \left| \int_{\mathcal{R}} f(r) \left(\sum_{k=1}^{K} \pi_{1,k} p^{1,k}_\Phi(r) \right) dr - \int_{\mathcal{R}} f(r) \left(\sum_{k=1}^{K} \pi_{0,k} p^{0,k}_\Phi(r) \right) dr \right|$$

2. *Linearity of the integral.* We now make explicit the contribution of each mixture component by using properties of the integral. Since integration is a linear operator and the mixture weights $\pi_{t,k}$ are constants with respect to the integration variable r, we have:

$$\int_{\mathcal{R}} f(r) \left(\sum_{k=1}^{K} \pi_{1,k} p^{1,k}_\Phi(r) \right) dr = \sum_{k=1}^{K} \pi_{1,k} \int_{\mathcal{R}} f(r)\, p^{1,k}_\Phi(r)\, dr,$$
$$\int_{\mathcal{R}} f(r) \left(\sum_{k=1}^{K} \pi_{0,k} p^{0,k}_\Phi(r) \right) dr = \sum_{k=1}^{K} \pi_{0,k} \int_{\mathcal{R}} f(r)\, p^{0,k}_\Phi(r)\, dr.$$

By substituting these expressions into the definition of the IPM then we can write:

$$\mathrm{IPM}_{\mathcal{G}}(p^1_\Phi, p^0_\Phi) = \sup_{f\in\mathcal{G}} \left| \sum_{k=1}^{K} \pi_{1,k} \int_{\mathcal{R}} f(r)\, p^{1,k}_\Phi(r)\, dr - \sum_{k=1}^{K} \pi_{0,k} \int_{\mathcal{R}} f(r)\, p^{0,k}_\Phi(r)\, dr \right|$$

3. *Expected value notation.* We now rewrite each integral in terms of expected value. Recall that, for any measurable function f and any density p on $\mathcal{R}$, $\mathbb{E}_p[f] := \int_{\mathcal{R}} f(r)\, p(r)\, dr$ Therefore, for each treatment group $t \in \{0,1\}$ and cluster $k \in \{1,\dots,K\}$, we have:

$$\int_{\mathcal{R}} f(r)\, p^{t,k}_\Phi(r)\, dr \;=\; \mathbb{E}_{p^{t,k}_\Phi}[f]$$

Using this notation, the IPM becomes:

$$\mathrm{IPM}_{\mathcal{G}}(p^1_\Phi, p^0_\Phi) = \sup_{f\in\mathcal{G}} \left| \sum_{k=1}^{K} \pi_{1,k}\, \mathbb{E}_{p^{1,k}_\Phi}[f] - \sum_{k=1}^{K} \pi_{0,k}\, \mathbb{E}_{p^{0,k}_\Phi}[f] \right|$$

4. *Isolation of local and structural contributions* At this point, the difference between treated and control mixtures combines two effects: *(i)* differences between the cluster-conditional distributions, i.e., discrepancies between treated and control units within the same cluster, $p^{1,k}_\Phi$ and $p^{0,k}_\Phi$, and *(ii)* differences between the mixture weights $\pi_{1,k}$ and $\pi_{0,k}$ (how much treated and controls end up in different clusters). To separate these contributions, we add and subtract the intermediate term $\sum_{k=1}^{K} \pi_{1,k}\, \mathbb{E}_{p^{0,k}_\Phi}[f]$ inside the absolute value:

$$\begin{aligned}\mathrm{IPM}_{\mathcal{G}}(p^1_\Phi, p^0_\Phi) &= \sup_{f\in\mathcal{G}} \left| \sum_{k=1}^{K} \pi_{1,k}\, \mathbb{E}_{p^{1,k}_\Phi}[f] - \sum_{k=1}^{K} \pi_{0,k}\, \mathbb{E}_{p^{0,k}_\Phi}[f] \right| \\ &= \sup_{f\in\mathcal{G}} \Bigg| \underbrace{\sum_{k=1}^{K} \pi_{1,k} \left(\mathbb{E}_{p^{1,k}_\Phi}[f] - \mathbb{E}_{p^{0,k}_\Phi}[f] \right)}_{\text{within-cluster (local) discrepancy}} + \underbrace{\sum_{k=1}^{K} (\pi_{1,k} - \pi_{0,k})\, \mathbb{E}_{p^{0,k}_\Phi}[f]}_{\text{across-cluster (structural) imbalance}} \Bigg|.\end{aligned}$$

5. *Triangle inequality and separation of contributions.* At this stage, the global discrepancy is expressed as the supremum of the sum of two terms: a *local* discrepancy within clusters and a *structural* imbalance across clusters. To upper bound this quantity, we apply the triangle inequality and basic properties of the supremum.

Let us define, for a fixed test function $f \in \mathcal{G}$,

$$A_f := \sum_{k=1}^{K} \pi_{1,k} \Big(\mathbb{E}_{p_\Phi^{1,k}}[f] - \mathbb{E}_{p_\Phi^{0,k}}[f] \Big), \qquad B_f := \sum_{k=1}^{K} (\pi_{1,k} - \pi_{0,k}) \, \mathbb{E}_{p_\Phi^{0,k}}[f].$$

A_f captures the discrepancy between treated and control distributions *within each cluster*, while B_f captures the discrepancy due to different cluster occupancies between the two groups. By the triangle inequality, for any $f \in \mathcal{G}$, $|A_f + B_f| \leq |A_f| + |B_f|$, and so:

$$\mathrm{IPM}_{\mathcal{G}}(p_\Phi^1, p_\Phi^0) = \sup_{f \in \mathcal{G}} |A_f + B_f| \;\leq\; \sup_{f \in \mathcal{G}} |A_f| + \sup_{f \in \mathcal{G}} |B_f|. \tag{12}$$

We now bound the two terms separately. Since the mixture weights satisfy $\pi_{1,k} \geq 0$ and $\sum_{k=1}^{K} \pi_{1,k} = 1$, we can use the subadditivity of the supremum to obtain

$$\begin{aligned} \sup_{f \in \mathcal{G}} |A_f| &= \sup_{f \in \mathcal{G}} \left| \sum_{k=1}^{K} \pi_{1,k} \Big(\mathbb{E}_{p_\Phi^{1,k}}[f] - \mathbb{E}_{p_\Phi^{0,k}}[f] \Big) \right| \\ &\leq \sum_{k=1}^{K} \pi_{1,k} \sup_{f \in \mathcal{G}} \left| \mathbb{E}_{p_\Phi^{1,k}}[f] - \mathbb{E}_{p_\Phi^{0,k}}[f] \right| \\ &= \sum_{k=1}^{K} \pi_{1,k} \; \mathrm{IPM}_{\mathcal{G}}\big(p_\Phi^{1,k}, p_\Phi^{0,k}\big). \end{aligned} \tag{13}$$

The second term is left in compact form and defines the residual *structural* discrepancy:

$$\delta(m) := \sup_{f \in \mathcal{G}} \left| \sum_{k=1}^{K} (\pi_{1,k} - \pi_{0,k}) \, \mathbb{E}_{p_\Phi^{0,k}}[f] \right|.$$

Combining equation 12 and equation 13, we finally obtain

$$\mathrm{IPM}_{\mathcal{G}}(p_\Phi^1, p_\Phi^0) \;\leq\; \sum_{k=1}^{K} \pi_{1,k} \; \mathrm{IPM}_{\mathcal{G}}\big(p_\Phi^{1,k}, p_\Phi^{0,k}\big) \;+\; \delta(m) \tag{14}$$

3. Conclusion of the Proof

By combining the IPM decomposition with the generalization bound for factual and counterfactual error Shalit et al. (2017), we obtain the final formulation of Theorem 3:

$$\mathrm{EMSE}_\Phi \leq 2 \left[\varepsilon_{F,1} + \varepsilon_{F,0} + B_\Phi \left(\sum_k \pi_k \mathrm{IPM}_{\mathcal{G}}(p_\Phi^{1,k}, p_\Phi^{0,k}) + \delta(m) \right) \right] \tag{15}$$

This proves that `HERMES` optimizes the error not through indiscriminate global alignment, but through a *structured local alignment* that preserves the essential causal distinctions of the latent space. The margin does not aim to eliminate global distributional differences, but to ensure that treated and control units are aligned only where their treatment effects are comparable. This transforms the CATE estimation problem from global *extrapolation* into structured local *interpolation*.

B Additional Results

B.1 Methodological Comparison of HERMES with Other Deep Learning Methods

Table 5 describes other state of the art CATE deep learning methods according to: *(i) Learning loss*, i.e., the core strategy for the loss function definition, *(ii) Dist. criterion*, i.e., criterion for the definition of treated/control distributions, *(iii) Balancing*, i.e., what is the basis for the balancing (global alignment only, neighbourhood-based alignment), *(iv) Latent space*, i.e., what notion of similarity defines "proximity", and *(v) Weakness*, i.e., the most critical weakness.

Table 5: Comparative analysis of CATE estimation by representation learning models

Model	Learning loss	Dist. criterion	Balancing	Latent space	Weakness
TARNet (Shalit et al., 2017)	*factual loss for outcome estimation*	*factual outcome*	*no explicit alignment between treated/control*	*unstructured*	*selection bias, counterfactual extrapolation*
CFRNet (Shalit et al., 2017)	*treated/control discrepancy penalty*	*treated/control distribution discrepancy*	*covariate similarity*	*covariate similarity based-structure*	*expensive loss computationally*
BCAUSS (Tesei et al., 2023)	*self-supervised balancing score*	*treated/control weighted covariate means discrepancy*	*global covariate*	*covariate similarity-based structure*	*learned weights quality-related performance*
CITE (Li & Yao, 2022)	*softmax-based contrastive loss*	*external, static propensity score*	*global covariate*	*propensity score similarity-based structure*	*static pairs, external model*
Our proposal	*SNN margin-based contrastive loss*	*internal, dynamic ITE predictions*	*global/local ITE*	*ITE similarity-based structure*	*warm-up phase; early pairs labeling noise*

B.2 Illustrative Example

To gain an intuitive understanding of how our method works, we provide an illustrative example. Let us consider the problem of choosing the optimal psychiatric drug, which represents a paradigmatic case where an accurate estimation of ITE is crucial. In mood disorders, many patients discontinue therapy prematurely or refuse to change molecules due to side effects or perceived low efficacy. Having a system that, based on clinical characteristics, indicates the treatment with the highest probability of benefit could improve therapeutic adherence and, consequently, health outcomes.

To isolate the methodological effects, a *simulated* dataset was created as follows. We remark that, the treatment and outcome generators defined to build such a dataset, have been calibrated to, from one hand, violate the *positivity* assumption, by creating regions in the covariate space where only one of the two treatments is (almost) observable, and from the other hand, introduce *confounding* to test the robustness. Each patient is described by four key *confounders* variables, i.e., which influences both the probability of prescription and the clinical outcome: *(i) depression severity* $X_1 \sim \texttt{TruncatedNormal}(\mu = 15, \sigma = 5, a = 0, b = 30)$; *(ii) age* $X_2 \sim \texttt{TruncatedNormal}(\mu = 45, \sigma = 25, a = 18, b = 85)$; *(iii) substance abuse* $X_3 \sim \texttt{Bernoulli}(p = 0.2 + 0.015 \cdot X_1)$; *(iv) access to care* $X_4 \in \{\texttt{low}, \texttt{high}\}$ with $\Pr(X_4 = \texttt{high} \mid X_2) = \sigma(-1.5 + 0.03 \cdot X_2)$, where $\sigma(\cdot)$ is the sigmoid function.

The probability of receiving the psychiatric drug ($T = 1$) is defined as follows:

$$\Pr(T = 1 \mid \mathbf{X}) = \sigma(-1.2 + 0.30 \cdot X_1 + 0.02 \cdot X_2 - 0.8 \cdot \mathbf{1}_{\{X_3=1\}} + 1.1 \cdot \mathbf{1}_{\{X_4=\texttt{high}\}})$$

so that more serious patients with better access to care are treated more frequently, while substance abuse reduces the likelihood of psychiatric drug prescriptions.

The binary outcome Y (*symptom improvement* within 6 months) is defined as follows:

$$\Pr(Y = 1 \mid T, \mathbf{X}) = \sigma\Big(-0.5 - 0.4 \cdot \frac{X_1}{30} + 0.03 \cdot X_2 - 0.8 \cdot \mathbf{1}_{\{X_3=1\}} + 1.0 \cdot \mathbf{1}_{\{X_4=\texttt{high}\}} + 0.6 \cdot T - 0.2 \cdot T \cdot \mathbf{1}_{\{X_3=1\}} \Big) \quad (16)$$

In this way, the drug *reduces* severity for most patients (+0.6 main effect) and the effect is attenuated in individuals with substance abuse (−0.2 interaction term). The dataset obtained, characterized by positivity violation and confounding, will be used to compare `HERMES` with baseline methods.

Figure 4 shows the covariates and outcome distributions in a *RCT* scenario, with 200,000 individuals of which 100,000 treated and 100,000 control. The distributions of *depression severity*, *age*, *substance abuse*,

and *access to care* are nearly identical between the *Drug* and *No Drug* groups. Randomization breaks any association between covariates and treatment assignment, allowing the difference in outcome (*symptom improvement*) to be interpreted as *causal effect* of drug.

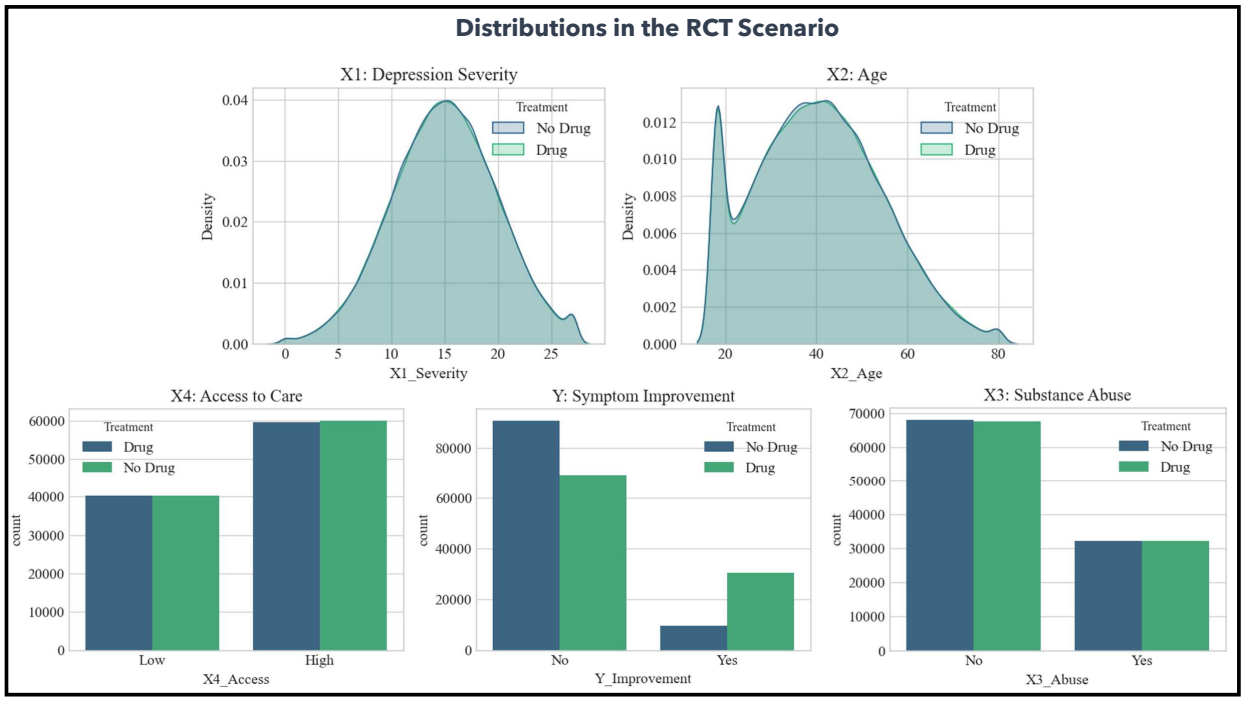

Figure 4: *RCT scenario*: comparison of covariates and outcome distributions.

Figure 5 shows the distributions in an *Observational* scenario (OBS), with 200,000 individuals of which 142,000 treated and 58,000 control. As we can see, patients receiving the drug tend to have more severe depression and higher average age; conversely, substance abuse and limited access to care are more frequent among the untreated. This pattern, known as *confounding by indication*, produces the paradox that the drug appears to worsen symptoms. In reality, the negative effect is due to the fact that the treatment is assigned to the most compromised patients from the beginning.

In the OBS scenario, to counteract confounding, the *propensity score* is often used. However, the function used to estimate such a probability affects the ability of the model to balance the distributions. Figure 6 shows two hypothetical functions analyzed in Tesei et al. (2023):

1. *BCE.* Binary cross-entropy "step" function that learns a threshold ($\approx$ 15 in Figure 6) on the *depression severity* covariate, above which $P(T{=}1 \mid X) = 1$, by correlating covariates to treatment assignment rather than the treatment effect; this produces extreme propensity scores, and highly dissimilar latent representations for treated versus control individual.

2. *BCAUSS.* Softer loss balancing function on the weighted representation of covariates: the probability of treatment varies continuously with the depression severity covariate, avoiding threshold-based decisions; the propensity scores are more overlapped on the covariates, indicating greater covariate balance in the latent space and more favorable conditions for estimating the causal effect (less residual confounding); by reducing distributional discrepancies between treated and control representations, i.e., enforcing $p(X \mid T{=}1) \approx p(X \mid T{=}0)$ for the learned representations, `BCAUSS` makes treatment assignment nearly independent of pre-treatment covariates conditional on X (approximately $T \perp (Y(0), Y(1)) \mid X$).

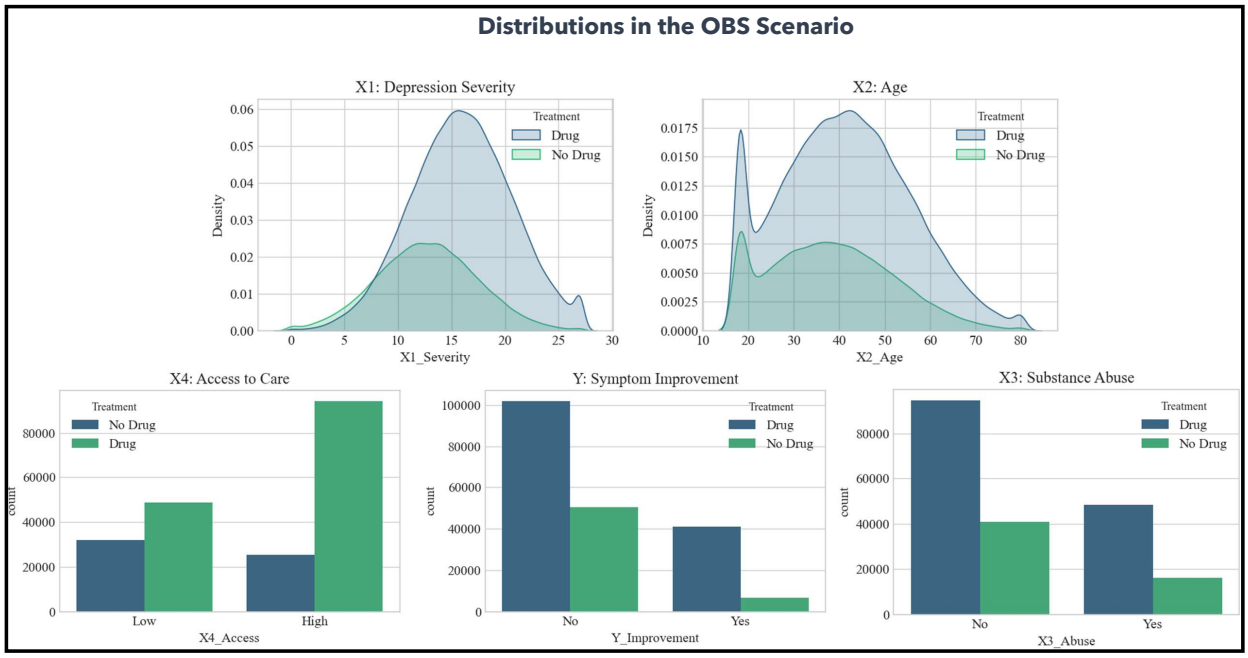

Figure 5: *OBS scenario*: comparison of covariates and outcome distributions.

Figure 6 shows propensity function and propensity score distributions under the two functions described above. For both such functions, the "left column" highlights the relationship between the *depression severity* covariate and the estimated treatment probability $P(T{=}1 \mid X)$, while the "right column" shows the histograms of the resulting propensity scores for the treated and control groups. As we can see on the "top", in *BCE*, the model learns a step function with a threshold at ~ 15, assigning treatment almost deterministically above the threshold and almost never below it. This yields extreme propensity scores (near 0 or 1) and little overlap between groups. On the "bottom", `BCAUSS` learns a balancing function over a weighted representation of the covariates, so treatment probability varies continuously with severity. The resulting propensity scores show substantially more overlap, indicating better covariate balance between treated and control individuals.

Our network extends the idea of balancing by introducing a *contrastive loss* that operates directly on the ITE. Each anchor sample is associated with: a *positive point* with a similar ITE; a *negative point* with a different ITE. As shown in Figure 7, the training encourages the network to map anchors and positives within a sphere of radius m, while keeping negatives outside this margin. By combining balancing and contrastive learning, our architecture allows for robust estimation of individual treatment effects even in the presence of strong confounding and positivity violations.

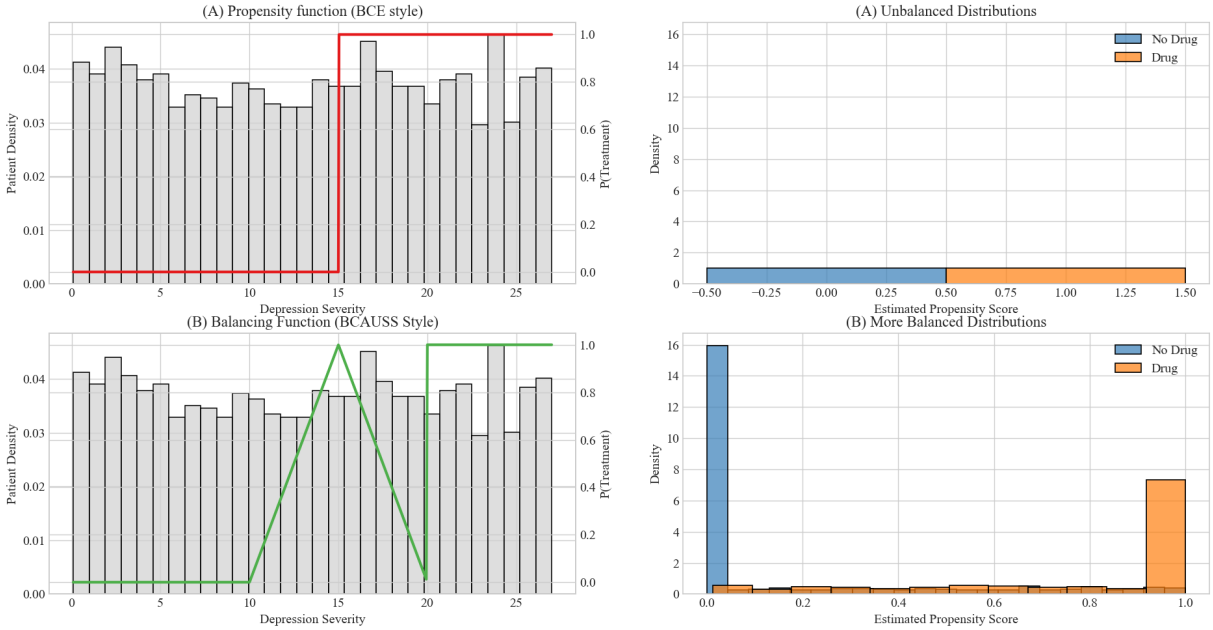

Figure 6: Propensity function and propensity–score distributions under two training objectives. *Left column:* relationship between the covariate *depression severity* and the estimated treatment probability $P(T{=}1 \mid X)$. *Right column:* histograms of the resulting propensity scores for the treated (*Drug*) and control (*No Drug*) groups. *Top row (BCE).* The model trained with Binary Cross-Entropy learns a near step function with a threshold at ~ 15, assigning treatment almost deterministically above the threshold and almost never below it. This yields extreme propensity scores (near 0 or 1) and little overlap between groups. *Bottom row (BCAUSS).* The BCAUSS objective learns a smoother balancing function over a weighted representation of the covariates, so treatment probability varies continuously with severity. The resulting propensity scores show substantially more overlap, indicating better covariate balance between treated and control units.

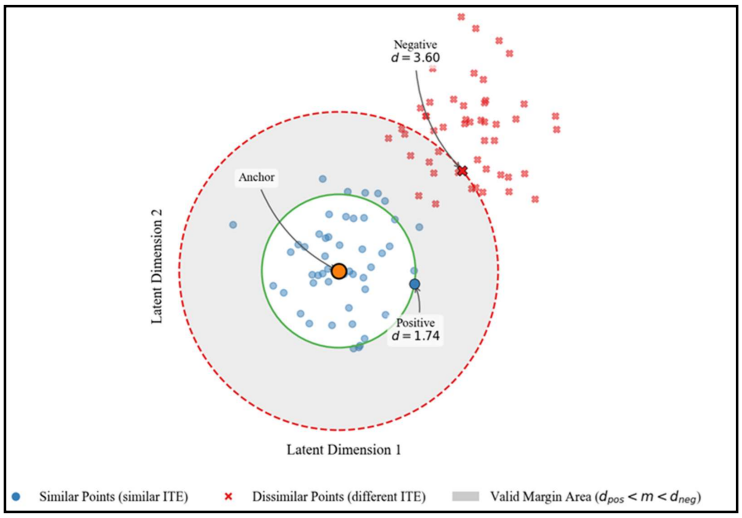

Figure 7: Impact of the *margin* m in the `HERMES`'s contrastive loss. Each **anchor** (solid circle) is paired with a **positive** with a similar ITE (blue circle) and a **negative** with a dissimilar ITE (red cross). The loss forces positive pairs to be *inside* a sphere of radius m and negative pairs *outside*.