

Improving Bias Correction Standards by Quantifying its Effects on Treatment Outcomes

Alexandre Abraham¹^[0000–0003–3693–0560] (✉) and Andrés Hoyos Idrobo²

¹ Implicity, Paris, France abraham.alexandre@gmail.com

² Rakuten Institute of Technology; Rakuten Group, Inc; Paris, France

Abstract. With the growing access to administrative health databases, retrospective studies have become crucial evidence for medical treatments. Yet, non-randomized studies frequently face selection biases, requiring mitigation strategies. Propensity score matching (PSM) addresses these biases by selecting comparable populations, allowing for analysis without further methodological constraints. However, PSM has several drawbacks. Different matching methods can produce significantly different Average Treatment Effects (ATE) for the same task, even when meeting all validation criteria. To prevent cherry-picking the best method, public authorities must involve field experts and engage in extensive discussions with researchers.

To address this issue, we introduce a novel metric, A2A, to reduce the number of valid matches. A2A constructs artificial matching tasks that mirror the original ones but with known outcomes, assessing each matching method’s performance comprehensively from propensity estimation to ATE estimation. When combined with Standardized Mean Difference, A2A enhances the precision of model selection, resulting in a reduction of up to 50% in ATE estimation errors across synthetic tasks and up to 90% in predicted ATE variability across both synthetic and real-world datasets. To our knowledge, A2A is the first metric capable of evaluating outcome correction accuracy using covariates not involved in selection. Computing A2A requires solving hundreds of PSMs, we therefore automate all manual steps of the PSM pipeline. We integrate PSM methods from Python and R, our automated pipeline, a new metric, and reproducible experiments into *popmatch*, our new Python package, to enhance reproducibility and accessibility to bias correction methods.

Keywords: Propensity score matching · Causal effect estimation · Python

1 Introduction

Propensity Score Matching (PSM) seeks to mitigate selection bias between control and treated populations by identifying comparable subpopulations based on potentially confounding covariates. According to the propensity score theorem [19], under strong ignorability – i.e., treatment assignment is independent of potential outcomes conditional on the confounding covariates – the propensity score can balance treatment groups and allow estimating causal effects. When applicable,

this method is convenient because it enables matching on a single value rather than full samples, as required in Mahalanobis matching, and it allows for the subsequent use of conventional statistical techniques. Since its initial definition three decades ago [19], PSM has undergone numerous methodological improvements [20], has been proven accurate by meta-analysis [16] and therefore recognised as a valid bias correction by French health authorities [5].

Despite these positive signals, PSM faces criticism for both theoretical and practical flaws [8]. The typical PSM pipeline [11], illustrated in Figure 1, includes three steps with multiple options (in blue) and two validation steps (in red). Guidelines recommend progressing as far as possible in the pipeline, making adjustments when validations fail, and backtracking when all alternatives are exhausted [11]. Unfortunately, the many options at each step lead to a combinatorial explosion of pipelines, each measuring a different Average Treatment Effect (ATE), with no established best practices for choosing the best model. Without strict validation, many results may seem valid, potentially misleading practitioners to select pipelines that fit their hypotheses rather than adhering to best practices.

Reporting difficulties. As each step of this pipeline is manual, the practitioner must carefully motivate his choices through standardized feature assessment before selection, evaluation of unmatched patients, visual inspection of propensity scores, and more [6,11]. This complexity hinders reproducibility and leads to significant shortcomings in method reporting. Systematic reviews of PSM applications reveal that many studies fail to adequately assess the balance of covariate after matching. This issue is observed in 17% of studies according to a 2008 study [1], 28% in 2011 [21], 26% in 2015 [13], 20% in 2017 [23], and in a substantial 48% in a 2020 study [4]. The situation is even more concerning when it comes to reporting the matching method itself with numbers as high as 30% in a 2008 study [1], 48% in 2011 [21], 33% in 2015 [13], and 30% in a 2017 study [23].

Efforts have been made to streamline these best practices, including the creation of visual aids in the MatchIt package [18,10], but ensuring the practitioner’s objectivity remains challenging. In this study, we lay the groundwork for a fully automated PSM system that, similar to autorank [7], could automatically make decisions and provide a comprehensive report detailing the rationale behind them.

High variability in valid ATE. The first validation step in the pipeline, specific to propensity-based methods, involves *visually inspecting* the propensity score distributions. The second step, common to all bias correction methods, is verifying covariate balance, usually based on Standardized Mean Difference (SMD). Model and hyperparameter choices for propensity estimation and matching can significantly impact variability in the results. This variability can stem from model dependence [8], unknown confounders, or inaccurate propensity estimation, all of which are difficult to diagnose. For example, MatchIt offers 11 propensity models and 8 matching methods, resulting in 88 possible matchings with varying ATE, many of them being valid according to our experiments. Stricter validation may help address this issue, which is the primary goal of this work.

While these issues are still debated, PSM remains a reference method in retrospective studies [23,4,22,15], necessitating a practical response to prevent decisions based on flawed interpretations. We tackle the issue of combinatorial explosion by enhancing method validation to minimize variance in estimated effects. Employing a fully automated pipeline, we investigate multiple PSM methods for each task, highlighting the variability in methods deemed valid by SMD. Then, we demonstrate that our metric A2A aids in reducing their quantity while retaining the most accurate ones. This work is organised as follows. We begin by highlighting the shortcomings of PSM assessment. Then, we present our evaluation task and how to create an artificial task from the data at hand and use it to validate a method. Finally, we expose our results and discuss them. Note that the code for all the presented methods and experiments is available in *popmatch*³, our Python package for PSM.

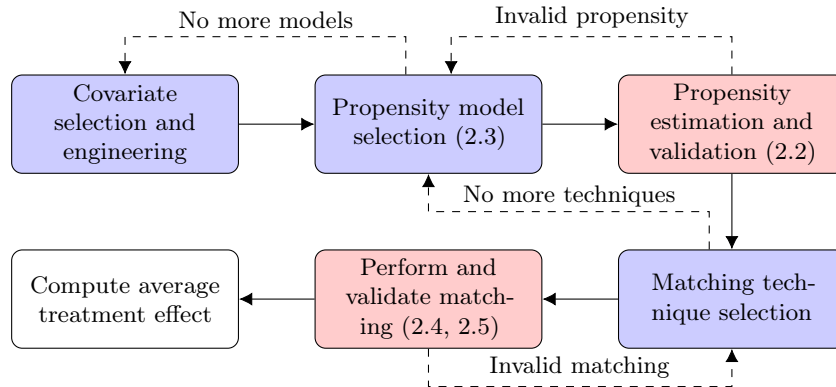


Fig. 1. Propensity score matching pipeline. Blue boxes indicate steps where the practitioner makes choices. Red boxes indicates steps ending with a validation. Backward arrows show points where the practitioner may revisit previous decisions.

2 Methods

PSM consists of finding, in a control and treated population, two exchangeable subsets for comparison. For generality, we refer to the two initial sets as X_0 and X_1 , and their matched counterparts as $\hat{X}_0 \subseteq X_0$ and $\hat{X}_1 \subseteq X_1$. The associated targets of interest are $Y_0, Y_1, \hat{Y}_0, \hat{Y}_1$. This work focuses on bipartite matching, where all patients from the smaller population are matched with patients from the larger one. This section outlines our automation of the PSM pipeline and the development of our metric, following the pipeline as presented in Figure 1. Instead of experimenting and backtracking like in a single experiment, we apply

³ Code is available at <https://github.com/NamelessAuthor/popmatch>

all available options for each task. This approach simulates the variability in results that would occur if different practitioners used different methods for the same study.

2.1 Real and synthetic tasks

Synthetic tasks simulate the treatment effect for each patient, allowing us to make an exact comparison between the measured and real ATE. For data generation, we draw inspiration from the setup B in [14] and their implementation in the CausalML Python package [2]. We use synthetic data comprising 3 000 samples with 10 numerical features and vary the number of confounders from 0 to 10 to observe their impact on the metrics.

Real-life tasks rely on openly available datasets that include a treatment variable, as summarised in Table 1. Missing data is imputed using the mean for continuous variables and the mode for categorical ones. The objective variable is binary for *Horse Colic* and *NHANES*, and real for *Groupon*. These datasets are used to analyse the behaviour of our metrics and not for performance evaluation since the real ATE is not known.

Table 1. Real-life dataset used for experiments. "Cont." represents continuous features, while "Cat." represents categorical ones.

Dataset	# Samples	# Cont.	# Cat.	Treatment	Prediction
Groupon ⁴	710	4	2	Discount applied	Revenue
Horse colic	300	7	13	Surgery or drug	Survival
NHANES ⁵	3974	4	15	Rheumatoid arthritis	Heart attack

2.2 Reminder: Computing propensity scores

The propensity score is the probability of receiving treatment given a set of covariates \mathbf{x} :

$$\widehat{\text{PS}}(\mathbf{x}) = \mathbb{P}[T = 1 \mid \mathbf{x}] + \epsilon, \quad (1)$$

where T is a binary treatment indicator, with $T = 1$ if the unit is treated and $T = 0$ otherwise, $\mathbf{x} = (x_1, x_2, \dots, x_n)$ is a vector of covariates, ϵ is the error term. Since the predicted probability and its complement can be used as denominators

⁴ Dataset provided by Harry Wang using Groupon data <https://www.kaggle.com/code/harrywang/propensity-score-matching-in-python/input>

⁵ Provided by Ehsan Karim using openly available data from the CDC <https://www.kaggle.com/code/wildscop/propensity-score-matching-on-nhanes>

in inverse propensity weighting, it is considered best practice to clip them within the interval $[0.05, 0.95]$ [11].

Propensity models are trained through feature preprocessing, feature selection, and model optimization. For feature preprocessing, continuous variables are centered and standardized to unit variance, while categorical variables are one-hot-encoded. Feature selection is unnecessary for our synthetic tasks as all features influence either propensity or outcome. Same goes for real tasks as the accompanying documentation suggests no feature is irrelevant.

We consider three propensity models that differ in nature and calibration methods. **Logistic Regression (LR)** is the most common and is calibrated using simple population weighting. **Chunked Logistic Regression (CLR)**, introduced in PsmPy [9], is a classic LR model using scikit-learn’s LR with a liblinear backend⁶. It is calibrated by segmenting the largest population into chunks the size of the smallest one and performing repeated predictions. This distinctive approach, not explained in their paper, yields results different from those of classic LR. We also use the more recent **Random Forest (RF)** [24], designed to handle non-linear data, and calibrate it using Platt’s scaling [17]. For all methods, we consider both raw predicted probabilities and logit link.

2.3 Automating propensity model hyperparameter selection

Propensity models may have hyperparameters, such as the number of trees for RF. Setting them based on prediction accuracy in a cross-validation (CV) scheme is impossible since propensity scores lack a ground truth. We propose to use a composite score to evaluate the best method. This score is computed on the left-out test set and is composed of:

Model performance as classification accuracy. Even if the true propensities are unknown, the binary assignment remains relevant information, as we expect the treated group to have a higher propensity:

$$\text{Accuracy} = \frac{1}{N_0} |\{\mathbf{x} \in X_0 : \widehat{\text{PS}}(\mathbf{x}) < 0.5\}| + \frac{1}{N_1} |\{\mathbf{x} \in X_1 : \widehat{\text{PS}}(\mathbf{x}) \geq 0.5\}|.$$

Extreme value ratio as the ratio of values outside of the interval $[0.05, 0.95]$ predicted in the left-out test set. While extreme propensities can be clipped, it is preferable to have none:

$$\text{Extremes} = \frac{1}{N} \left[|\{\mathbf{x} \in (X_0 \cup X_1) : \widehat{\text{PS}}(\mathbf{x}) \notin [0.05, 0.95]\}| \right].$$

Overlap coefficient [12] between the normalized histograms of the valid propensity scores of both populations in strata of 0.1. This automates the *visual inspection* recommended in best practices [11].

$$\text{NH}_i(X) = \frac{1}{|X|} |\{\mathbf{x} \in X : \widehat{\text{PS}}(\mathbf{x}) \in [i \times 0.1 - 0.05, i \times 0.1 + 0.05]\}|$$

⁶ Surprisingly, using a different backend degrades CLR’s performance. This behavior warrants further investigation.

$$\text{Overlap} = \sum_{i \in [1, \dots, 9]} \min(\text{NH}_i(X_0), \text{NH}_i(X_1))$$

This idea is reminiscent of stratified propensity matching and is depicted in Figure 2. Note that throughout the paper, a propensity method is deemed invalid if this overlap falls below 0.5.




Fig. 2. Example of propensity score histograms for control and treated population. The hatched area corresponds to the overlap between the two.

We use the following composite score in a 5-fold cross-validation (CV) to select the best models:

$$\text{Score} = \text{Accuracy} + (1 - \text{Extremes ratio}) + \text{Overlap ratio}.$$

2.4 Validating matching with covariate balancing

SMD computes the distance between two distribution of features, typically control and treated populations. SMD for continuous variables is usually computed using Cohen’s D with pooled standard deviation, *i.e.* the T-test effect size. For categorical variables, we use Cramér’s V, the χ^2 effect size, as there is no standard:

$$\text{Cohen's D} = \frac{\mu_0 - \mu_1}{\sqrt{\frac{(N_0-1)\cdot\sigma_0^2 + (N_1-1)\cdot\sigma_1^2}{N_0 + N_1 - 2}}}, \quad \text{Cramér's V} = \sqrt{\frac{\chi^2}{N \cdot \min(c-1, r-1)'}}$$

with μ_0 (resp. μ_1) the means for the first (resp. second) populations, N_0 (resp. N_1) their respective cardinalities, σ_0 (resp. σ_1) their respective standard deviations, N the total number of observations, and c (resp. r) the number of columns (resp. rows) in the contingency table. French Health Authorities consider that having an SMD below 10% after correction validates the bias correction method. We use the same criterion in our experiments.

2.5 Contribution: Validating ATE and matching with a new metric

ATE is simply the difference between the average effect in both populations.

$$\text{ATE}(Y_0, Y_1) = \mathbb{E}[Y_1] - \mathbb{E}[Y_0]. \quad (2)$$

The ability of a PSM to retrieve the true ATE can only be assessed using synthetic tasks or through literature on real datasets [16]. There exists no method to select the best technique for the specific task being studied. We propose achieving this by generating artificial matching tasks using the studied data. These tasks replicate biases similar to those in the original task but with known ATE. We posit that an approach’s ability to recover the known ATE in these artificial tasks measures its effectiveness in retrieving the ATE in the original task. This procedure evaluates the entire matching process, including the interaction between propensity estimation and matching. We term this new metric *A2A*.

Note that in this manuscript, we make a clear distinction between two types of synthetic problems: *artificial tasks*, which are derived from a homogeneous population and possess a null ATE, used for A2A computation, and *synthetic tasks*, where we simulate an entire matching problem with a known ATE to evaluate matching methods. We will maintain this terminology throughout the rest of the manuscript to differentiate between these two types of scenarios.

Artificial task from real world data The core idea of our work is to establish a matching task with a predetermined outcome. We suggest partitioning a sub-population uniformly, for example controls X_0 , into two subsets $X_0^{(0)}$ and $X_0^{(1)}$. As both subsets originate from the control population, we can ascertain that the treatment effect remains consistent across both, leading to an estimated ATE of zero post PSM.

Selecting subsets. To assess the PSM method’s efficacy in addressing the reference task, we design our artificial tasks to mimic the biases present in the reference scenario. We select subsets and ensure their unadjusted ATE matches that of the reference task, as shown in Equation 4a, while also enforcing equivalent initial SMD conditions, as illustrated in Equation 4b. Since the artificial task uses half the dataset, we introduce only half the biases to avoid making it too challenging and ensure a good spread of scores. This adjustment doesn’t impact the final method rankings. Lastly, we determine the cardinality of the two subsets to be proportional to that of the reference task, as indicated in Equation 3b.

Thus, selecting subsets to match boils down to solving the following optimization problem:

$$\{X_0^{(0)}, X_0^{(1)}\} \leftarrow \arg \min_{\{X_0^{(0)}, X_0^{(1)}\}} \mathcal{L}(X, Y, \{X_0^{(0)}, X_0^{(1)}\}), \quad (3a)$$

$$\text{s.t. } X_0^{(0)}, X_0^{(1)} \subseteq X_0, \quad X_0^{(0)} \cap X_0^{(1)} = \emptyset, \quad \underbrace{\frac{|X_0|}{|X_1|} = \frac{|X_0^{(0)}|}{|X_0^{(1)}|}}_{\text{Same size ratio}}. \quad (3b)$$

where,

$$\mathcal{L}\left(X, Y, \{X_0^{(0)}, X_0^{(1)}\}\right) = \left(\frac{1}{2} \underbrace{\text{ATE}(Y_0, Y_1)}_{\text{reference task}} - \underbrace{\text{ATE}\left(Y_0^{(0)}, Y_0^{(1)}\right)}_{\text{artificial task}}\right)^2 \quad (4a)$$

$$+ \left(\frac{1}{2} \underbrace{\text{SMD}(X_0, X_1)}_{\text{reference task}} - \underbrace{\text{SMD}\left(X_0^{(0)}, X_0^{(1)}\right)}_{\text{artificial task}}\right)^2. \quad (4b)$$

Where $X_0^{(0)}$, resp. $X_0^{(1)}$, represents our artificial control, resp. treated, population. Without losing generality, we assume $|X_0| \geq |X_1|$. Ultimately, we obtain a task similar to the original but with a known ATE of 0.

This set building task is easier to solve than PSM and there as it allows for the free swapping of samples between populations. Thus it does not require sophisticated methods. Due to the absence of a well-defined gradient in the optimization problem, we use a hill-climbing method for minimization, as detailed in Algorithm A.1. In our experiments, the algorithm converged within seconds, making a more complex approach, *i.e.*, simulated annealing, unnecessary. Also, this method ensures non-overlapping subsets by construction and injects diversity into our artificial tasks as it reaches local minimums, while other methods aim for the global one.

A2A: A new metric to assess the best performing PSM pipeline. Identifying the best method involves minimizing the difference in ATE measured on the matched populations, $\hat{Y}_0^{(0)}$ and $\hat{Y}_1^{(0)}$:

$$\text{A2A} := \left| \text{ATE}\left(\hat{Y}_0^{(0)}, \hat{Y}_0^{(1)}\right) \right|. \quad (5)$$

We call this metric A-to-A, or A2A for short, reminiscently of the A/B test wording, since both populations used in the task come from the A population. The closer it is to zero, the better. To obtain the A2A score, we compute 100 bootstraps of the entire A2A process, from task creation to ATE estimation, on the larger population, whether control or treated, and take the mean of the resulting values. Unlike SMD, A2A is a relative metric: Lower values indicate better performance, but there is no guarantee that the debiasing problem is solvable, making it impossible to set a fixed threshold like SMD.

Combining SMD and A2A Since SMD is an absolute value, it allows for the use of a fixed threshold for method selection. In contrast, A2A is inherently relative and cannot use a fixed threshold. In order to combine both metrics in a single method selection process, we tested three strategies. The most straightforward, Min A2A, selects the best performing matching in terms of A2A among those deemed valid by SMD. For fairness, we added its counterpart Min SMD. The

other two strategies are more permissive, selecting a set of valid PSMs similar to SMD thresholding. SMD \times A2A relies on an algorithmic approach, applying DBScan [3] to cluster methods based on their SMD and A2A scores. Selecting the best methods boils down to keeping the cluster that contains the PSM with the lowest A2A score. The Pareto method, stemming from multi-objective optimization, selects the Pareto optimal PSM according to both A2A and SMD, among PSM deemed valid by SMD.

2.6 Contribution: Popmatch python package

The most commonly used implementation of PSM algorithms is the R package MatchIt [20]. While a Python counterpart exists, PsmPy [9], it includes only a single matching method. To improve reproducibility and broaden accessibility for Python-based data scientists, we introduce *popmatch*, a package that includes the original implementation of PsmPy and a Python interface to the methods offered by MatchIt. The code for downloading tasks, performing PSM selection, method evaluation, and all the experiments in this paper is available and reproducible.

2.7 Experimental setup

For each task, as required for A2A computation, we create 100 bootstrapped artificial tasks. We also use these tasks to compute metrics and gain insights, as described in this section. Since our datasets are tabular and our algorithms are easily tractable, all experiments can be completed within a few hours.

Because MatchIt’s propensity estimation and matching are tightly linked, we did not separate them. We retained methods that ran without errors on all bootstraps and are closest to 1-to-1 matching, specifically *nearest* and *optimal*, along with propensity score estimation methods *GLM*, *GAM*, *Elasticnet*, *Rpart*, *CBPS*, and *Bart*⁷. On the Python side, we combined LR, CLR, and RF propensity estimation methods with PsmPy’s matching.

3 Results

3.1 Validating propensity scores

As explained in Section 2.3, propensity estimation is considered valid if the overlap between the two distributions is at least 50%. Table 2 shows these ratios and indicates whether each method meets the threshold.

Regarding model validity, we first observe that applying the logit link, a recommended measure enabled by default in PsmPy, tends to yield invalid propensity models. Second, PsmPy’s CLR achieves remarkable results compared to LR and is the only method able to yield valid models on the Horse and NHANES datasets. Although this is clearly the optimal choice, we keep all models without logit link for the rest of our analysis.

⁷ The MatchIt package refers to “distance” as a method to estimate propensity scores.

Table 2. Overlap ratio of propensity scores depending on models and datasets. Colored cells indicate scores below 0.5. Bold numbers indicate the largest overlap row-wise.

Transform		None			Logit		
Model	LR	RF	CLR	LR	RF	CLR	
Groupon	0.75	0.70	0.67	0.25	0.27	0.26	
Horse	0.51	0.29	0.84	0.18	0.28	0.19	
NHANES	0.47	0.46	0.83	0.12	0.00	0.19	
Synth. data							
# confounders	0	1.00	0.86	0.97	0.00	0.62	0.45
	1	1.00	0.81	0.97	0.00	0.61	0.43
	2	0.85	0.60	0.97	0.36	0.41	0.44
	3	0.99	0.73	0.96	0.20	0.53	0.43
	4	1.00	0.87	0.97	0.17	0.52	0.43
	5	0.88	0.91	0.95	0.38	0.57	0.44
	6	1.00	0.86	0.97	0.00	0.27	0.40
	7	1.00	0.89	0.97	0.00	0.59	0.44
	8	0.88	0.83	0.97	0.35	0.40	0.43
	9	0.91	0.96	0.98	0.38	0.76	0.43
	10	0.91	0.90	0.96	0.39	0.04	0.44

3.2 Assessing SMD as a performance evaluator

Having artificial tasks with known ATE provides an opportunity to evaluate SMD’s ability to predict the performance of matching methods. We assess this based on the correlation between SMD’s ranking and the ground truth. Additionally, we assess the correlation with the magnitude of the correction applied, *i.e.*, the difference between the uncorrected ATE and the corrected one. We use Kendall’s τ correlation score due to our limited assumptions about the predicted ATE. As a sanity check, we compare it to a random ranking.

Table 3 shows that, as expected, the SMD ranking does not correlate with a random one. Most correlations between SMD and the ground truth are between 0.1 and 0.2, indicating a weak correlation. However, SMD demonstrates a stronger negative correlation with magnitude, revolving around -0.3 across most tasks. As lower SMD values are better, this implies that SMD measures the extent of correction rather than the accuracy or unbiasedness of the final result. This behavior aligns with the expectation that feature differences decrease with correction. However, the challenge is finding the optimal balance, as overcorrection can occur, leading to inaccuracies in the estimated effect.

3.3 Validity of matching methods

Table 4 presents results from real datasets, illustrating the trade-offs between both metrics. Methods deemed invalid by SMD are marked in red. Our analysis underscores the need for our approach: no single method consistently outperforms all others, making model selection essential.

Table 3. Ranking based correlation between SMD and the amount of correction applied, the ground truth, and random values as sanity check.

Dataset	Correction magnitude		Ground truth		Random		
	Kendall's τ	pValue	Kendall's τ	pValue	Kendall's τ	pValue	
Groupon	-0.30 ± 0.58	0.02	0.20 ± 0.45	0.05	-0.00 ± 0.16	0.51	
Horse	0.05 ± 0.30	0.34	0.05 ± 0.26	0.39	0.00 ± 0.15	0.56	
NHANES	-0.01 ± 0.56	0.11	0.39 ± 0.28	0.16	0.01 ± 0.17	0.53	
Synth. data							
# confounders	0	-0.39 ± 0.27	0.16	0.17 ± 0.31	0.31	0.01 ± 0.18	0.48
	1	-0.35 ± 0.27	0.14	0.16 ± 0.26	0.33	0.01 ± 0.17	0.52
	2	-0.39 ± 0.22	0.11	0.18 ± 0.23	0.32	0.01 ± 0.17	0.50
	3	-0.40 ± 0.21	0.11	0.20 ± 0.26	0.34	-0.02 ± 0.17	0.49
	4	-0.31 ± 0.31	0.14	0.15 ± 0.23	0.39	-0.00 ± 0.17	0.48
	5	-0.35 ± 0.26	0.14	0.14 ± 0.20	0.41	0.01 ± 0.19	0.44
	6	-0.40 ± 0.26	0.11	0.19 ± 0.23	0.31	0.02 ± 0.18	0.49
	7	-0.30 ± 0.31	0.22	0.12 ± 0.24	0.36	-0.02 ± 0.16	0.47
	8	-0.39 ± 0.31	0.15	0.11 ± 0.19	0.41	-0.01 ± 0.17	0.50
	9	-0.32 ± 0.27	0.14	0.12 ± 0.23	0.33	0.02 ± 0.18	0.45
	10	-0.30 ± 0.30	0.18	0.08 ± 0.24	0.38	-0.01 ± 0.17	0.49

We also observe that the metrics are complementary. For instance, ElasticNet and Rpart are identified as the best by A2A for Groupon but are invalid by SMD, while GAM and GLM show a similar pattern for Horse. Conversely, GAM and GLM are rated best by SMD for Groupon and NHANES but are not top-rated according to A2A. This trade-off is illustrated in Figure B.3.

Regarding matching methods, Bipartify stands out as the only method providing valid matches according to SMD for all datasets, although only one model is selected as the best for NHANES. Overall, PsmPy's CLR offers the best trade-off between metrics, despite narrowly failing to have a valid model selected for Horse.

3.4 Combining SMD and A2A: A better matching method evaluation

As demonstrated, model dependence can result in significant variance in the estimations of models deemed valid. We report this variance as the range between the highest and lowest ATE estimations among valid PSM methods, with a smaller range being preferable. Since Min A2A and Min SMD select only one model each, their range is null. Secondly, we measure the average error of validated models on synthetic tasks with known ground truth.

Table 5 demonstrates that SMD×A2A and Min A2A, which prioritize A2A, lead in terms of range and error when the number of confounders is low. As the number of confounders increases, Pareto becomes the best among methods selecting multiple PSMs, while Min SMD consistently outperforms the others.

A significant challenge in bias correction is its one-shot nature, where there is no room for error. Thus, a method's reliability is judged by its weakest

Table 4. SMD and A2A for valid matching methods on real datasets. Red cells indicate results considered invalid according to SMD.

Matching	SMD			A2A		
	Groupon	Horse	NHANES	Groupon	Horse	NHANES
<i>Optimal</i>						
ElasticNet	0.118	0.128	0.046	84.327	0.036	0.002
Rpart	0.164	0.151	0.070	88.880	0.038	0.003
CBPS	0.052	0.117	0.016	118.526	0.036	0.006
Bart	0.034	0.106	0.024	145.374	0.053	0.003
GAM	0.026	0.115	0.015	220.050	0.035	0.003
GLM	0.026	0.115	0.015	220.050	0.035	0.003
<i>Bipartify</i>						
RF	0.051	0.077	0.044	588.262	0.043	0.002
LR	0.039	0.086	0.032	624.086	0.048	0.001
CLR	0.054	0.076	0.056	834.433	0.047	0.002
<i>PsmPy</i>						
CLR	0.026	0.107	0.028	292.465	0.037	0.004
RF	0.051	0.179	0.166	1491.992	0.039	0.001
LR	0.043	0.158	0.165	1597.057	0.041	0.003

performance. Both Min A2A and Min SMD excel in their respective domains, but they falter significantly on Synthetic 3 (Min SMD) and Synthetic 6 (Min A2A). Overall, Pareto strikes the best balance, making it the most reliable PSM selection method.

On real datasets, Pareto achieves the best range for Groupon and slightly improves over SMD for NHANES, while SMD×A2A excels for Horse and NHANES. Based on these observations, we hypothesize that Groupon likely has more confounders.

4 Discussion

Applying PSM is not inherently challenging for a trained practitioner. However, the current SMD-based validation is too lenient, allowing multiple methods with varying ATEs to be considered valid. This situation forces practitioners to narrow their method search and justify decisions to authorities, who must understand and evaluate the study, potentially challenging choices due to the absence of established best practices for selecting propensity computation methods. Furthermore, without rigorous reporting standards, errors or data manipulation can go unnoticed, potentially leading to poor decisions.

A2A effectively reduces the variance in estimated ATEs without compromising accuracy. We have observed that SMD excels at estimating the amount of bias correction applied and measuring ATE, especially in the presence of confounders,

Table 5. Metric-based performance of PSM on real and artificial tasks. *Range of ATE values* represents the difference between the maximum and minimum ATE across valid methods. For synthetic tasks with known ATE, the mean squared error is reported per method. Since Min A2A and SMD select only one PSM, their range is 0 and not shown.

Dataset	Range of ATE values			ATE estimation error					
	SMD	SMD ×A2A	Par- eto	SMD	SMD ×A2A	Par- eto	Min A2A	Min SMD	
Groupon	1417	1176	659	<i>Real ATE unknown</i>					
Horse	0.014	0.000	0.014						
NHANES	0.022	0.000	0.005						
Synth. data									
# confounders	0	0.067	0.032	0.031	0.054	0.022	0.035	0.063	0.006
	1	0.099	0.003	0.003	0.070	0.043	0.045	0.039	0.052
	2	0.099	0.013	0.013	0.063	0.043	0.046	0.030	0.058
	3	0.122	0.079	0.094	0.069	0.062	0.066	0.028	0.103
	4	0.096	0.046	0.012	0.059	0.060	0.046	0.041	0.039
	5	0.087	0.064	0.018	0.057	0.057	0.026	0.013	0.039
	6	0.095	0.000	0.028	0.066	0.094	0.063	0.094	0.049
	7	0.067	0.043	0.033	0.041	0.046	0.037	0.074	0.030
	8	0.046	0.029	0.007	0.034	0.034	0.030	0.035	0.025
	9	0.079	0.070	0.038	0.049	0.053	0.043	0.093	0.019
	10	0.078	0.038	0.032	0.052	0.051	0.030	0.023	0.044

where propensity and outcome are explained by the same covariates. However, its weaknesses emerge when the factors explaining the outcome and the selection are disjoint, particularly when the number of confounders is low. In these cases, A2A proves its usefulness through its measurement of ATE on artificial tasks. This difference makes sense if we consider that propensity-based correction are only able to accurately balance covariates that are linked to selection. In the presence of confounders, correcting propensity often results in well-corrected ATE. However, when ATE is influenced by covariates unrelated to propensity, A2A becomes crucial in selecting models that provide the most accurate ATE correction. This explains A2A’s relative nature. While each method excels in its domain, combining them in a Pareto optimal approach ensures a stable model selection method, effective regardless of the number of confounders.

The best PSM pipeline. While this is not the primary objective of our study, we conducted comparisons of matching methods on our tasks. The main finding is that applying logit on predicted probabilities, or using the emphnearest matching in MatchIt, results in models that are considered invalid based on the SMD. PsmPy [9] demonstrated superior performance compared to MatchIt [20] without identifying a fundamental cause. Our study, which integrates all methods into a single package, reveals that this may be due to CLR, PsmPy’s block-based model calibration. CLR shows significantly better overlap in propensities (Table 2) and

improved metric performance on real datasets, particularly Groupon (Table 4). This motivates the integration of CLR into MatchIt for further experiments.

Towards an Automated Pipeline and Benchmark for PSM. This work introduces an automated pipeline capable of performing PSM that satisfies all the requirements of the French Health Authorities. Such a standardized automated baseline could provide a valuable foundation for expert discussions between authorities and applicants. Running it serves as a proof of good faith, and health authorities may request it when further validation is required. This also paves the way for the development of PSM comprehensive benchmarks.

Conclusion. Propensity-based methods can compensate for selection bias based on given covariates. We have shown that SMD effectively evaluates these methods for this specific purpose. Incidentally, they also provide good compensation for causal effect estimation when the covariates involved in outcome estimation are the same as those in selection, *i.e.*, they are confounders. However, when covariates are not confounders or in the presence of mixed effects, SMD may fail to select the best method, and propensity correction might be inadequate.

This study introduces A2A, a new metric that offers insight into bias correction by evaluating the entire matching pipeline, including target effect estimation, on the data at hand. A2A evaluates a correction method based on its outcome correction capability, which, to our knowledge, is a novel approach. In scenarios with a low number of confounders, A2A is superior to SMD for selecting the best matching method. Since the number of confounders can be hypothesized but not definitively determined, it is challenging to identify the most accurate metric. To address this, we propose combining them with a Pareto optimal selection to leverage the strengths of both metrics and minimize estimation error.

Limitations. Although promising, our study would benefit from replication on a broader range of datasets to generate more insights and improve the existing pipeline, in particular by replacing DBScan by a simpler and more interpretable method. Lastly, this study lays the groundwork for future research. If this analysis is limited to PSM, the same strategy can be applied on other bias correction techniques such as Inverse Propensity Weighting. Additionally, A2A currently functions as a relative metric for refining PSM selection, necessitating the computation of all pipelines to identify the best ones. Our goal is to develop an absolute metric usable with a threshold, similar to SMD. Furthermore, the scheme behind A2A could potentially be adapted to find weights, complementary to propensity, that allows for a better correction of the ATE.

Acknowledgments. This research was conducted on our spare time. Alexandre Abraham owes his deepest gratitude to his family, Hinata and Audrey for providing him with the time to pursue this work. He also thanks Implicity for indirectly contributing through the experiment that inspired this work.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Austin, P.C.: A critical appraisal of propensity-score matching in the medical literature between 1996 and 2003. *Statistics in medicine* **27**(12), 2037–2049 (2008)
2. Chen, H., Harinen, T., Lee, J.Y., Yung, M., Zhao, Z.: Causalml: Python package for causal machine learning. arXiv preprint arXiv:2002.11631 (2020)
3. Ester, M., Kriegel, H.P., Sander, J., Xu, X., et al.: A density-based algorithm for discovering clusters in large spatial databases with noise. In: *kdd*. vol. 96.34, pp. 226–231 (1996)
4. Grose, E., Wilson, S., Barkun, J., Bertens, K., Martel, G., Balaa, F., Abou Khalil, J.: Use of propensity score methodology in contemporary high-impact surgical literature. *Journal of the American College of Surgeons* **230**(1), 101–112 (2020)
5. HAS: Méthodologie pour le développement clinique des dispositifs médicaux. Tech. rep., Haute autorité de Santé (2021)
6. Heinrich, C., Maffioli, A., Vazquez, G.: A primer for applying propensity-score matching. *Impact-Evaluation Guidelines* (2010)
7. Herbold, S.: Autorank: A python package for automated ranking of classifiers. *Journal of Open Source Software* **5**(48), 2173 (2020). <https://doi.org/10.21105/joss.02173>, <https://doi.org/10.21105/joss.02173>
8. King, G., Nielsen, R.: Why propensity scores should not be used for matching. *Political analysis* **27**(4), 435–454 (2019)
9. Kline, A., Luo, Y.: Psmpy: a package for retrospective cohort matching in python. In: 2022 44th Annual International Conference of the IEEE Engineering in Medicine & Biology Society (EMBC). pp. 1354–1357. IEEE (2022)
10. Lee, J., Little, T.D.: A practical guide to propensity score analysis for applied clinical research. *Behaviour Research and Therapy* **98**, 76–90 (2017)
11. Li, M.: Using the propensity score method to estimate causal effects: A review and practical guide. *Organizational Research Methods* **16**(2), 188–226 (2013)
12. McGill, M.: An evaluation of factors affecting document ranking by information retrieval systems. ERIC (1979)
13. McMurry, T.L., Hu, Y., Blackstone, E.H., Kozower, B.D.: Propensity scores: methods, considerations, and applications in the journal of thoracic and cardiovascular surgery. *The Journal of thoracic and cardiovascular surgery* **150**(1), 14–19 (2015)
14. Nie, X., Wager, S.: Quasi-oracle estimation of heterogeneous treatment effects. *Biometrika* **108**(2), 299–319 (2021)
15. Nowak, M.M., Niemczyk, M., Florczyk, M., Kurzyna, M., Pączek, L.: Effect of statins on all-cause mortality in adults: A systematic review and meta-analysis of propensity score-matched studies. *Journal of Clinical Medicine* **11**(19), 5643 (2022)
16. Olmos, A., Govindasamy, P.: Randomized experiments vs. propensity scores matching: A meta-analysis (2014)
17. Platt, J., et al.: Probabilistic outputs for support vector machines and comparisons to regularized likelihood methods. *Advances in large margin classifiers* **10**(3), 61–74 (1999)
18. Randolph, J.J., Falbe, K.: A step-by-step guide to propensity score matching in r. *Practical Assessment, Research & Evaluation* **19** (2014)
19. Rosenbaum, P.R., Rubin, D.B.: The central role of the propensity score in observational studies for causal effects. *Biometrika* **70**(1), 41–55 (1983)
20. Stuart, E.A., King, G., Imai, K., Ho, D.: Matchit: nonparametric preprocessing for parametric causal inference. *Journal of statistical software* (2011)

21. Thoemmes, F.J., Kim, E.S.: A systematic review of propensity score methods in the social sciences. *Multivariate behavioral research* **46**(1), 90–118 (2011)
22. Wang, J.: To use or not to use propensity score matching? *Pharmaceutical Statistics* **20**(1), 15–24 (2021)
23. Yao, X.I., Wang, X., Speicher, P.J., Hwang, E.S., Cheng, P., Harpole, D.H., Berry, M.F., Schrag, D., Pang, H.H.: Reporting and guidelines in propensity score analysis: a systematic review of cancer and cancer surgical studies. *JNCI: Journal of the National Cancer Institute* **109**(8), djw323 (2017)
24. Zhao, P., Su, X., Ge, T., Fan, J.: Propensity score and proximity matching using random forest. *Contemporary clinical trials* **47**, 85–92 (2016)

A Algorithms

Let \mathbf{x}_j be the j th row of X .

Data: Data $(\{X_0, X_1\}, Y)$, cluster membership probabilities $p \in [0, 1]^2$, loss function \mathcal{L} , number of iterations K .

$\mathbf{c}_j \leftarrow$ Assign $\mathbf{x}_j \in X_0$ to a cluster at random following $p, \forall j \in \{0, \dots, |X_0|\}$;

$Z^h \leftarrow \{\mathbf{x}_j \in X_0 : \mathbf{c}_j = h, \forall j \in \{0, \dots, |X_0|\}\}$, where $h = \{0, 1\}$;

$l^{\text{current}} \leftarrow \mathcal{L}(X, Y, \{Z^{(0)}, Z^{(1)}\})$ (evaluate using Eq. 4);

for $i \leftarrow 1$ **to** K **do**

$j, k \leftarrow$ Draw $(j, k) \in \{0, \dots, |X_0|\}^2$ at random with $\mathbf{c}_j \neq \mathbf{c}_k$

$\mathbf{c}_j, \mathbf{c}_k \leftarrow \mathbf{c}_k, \mathbf{c}_j$ (swap cluster assignment)

$Z^h \leftarrow \{\mathbf{x}_j \in X_0 : \mathbf{c}_j = h, \forall j \in \{0, \dots, |X_0|\}\}$, where $h = \{0, 1\}$;

$l^{\text{new}} \leftarrow \mathcal{L}(X, Y, \{Z^{(0)}, Z^{(1)}\})$

if $l^{\text{new}} < l^{\text{current}}$ **then**

$l^{\text{current}} \leftarrow l^{\text{new}}$

else

$\mathbf{c}_j, \mathbf{c}_k \leftarrow \mathbf{c}_k, \mathbf{c}_j$ (swap cluster assignment)

end

end

return $\{Z_0^{(0)}, Z_0^{(1)}\}$ (selected subsets/clusters);

Algorithm A.1: Selecting subsets to match by Hill-Climbing clustering

B SMD vs. A2A

This figure illustrates the trade-off between the two metrics. It's challenging to determine the superior method among Bipartify LR, and GLM/GAM solely from this plot. Hence, we opt for the Pareto optimal selection of methods. PsmPy's RF could potentially be considered Pareto-optimal, but it is effectively eliminated because its SMD exceeds 10




Fig. B.3. Values for SMD and A2A for the task NHANES. Methods above 0.1 in SMD are invalid. Bipartify LR and GLM form the Pareto optimum according to both metrics.