

MultiMBNN: Matched and Balanced Causal Inference with Neural Networks

Ankit Sharma¹, Garima Gupta¹, Ranjitha Prasad², Arnab Chatterjee¹,
Lovekesh Vig¹ and Gautam Shroff¹

1- TCS Research, Delhi, India

2- Indraprastha Institute of Information Technology, Delhi, India

Abstract. Causal inference (CI) in observational studies has received a lot of attention in healthcare, education, ad attribution, policy evaluation, etc. Confounding is a typical hazard, where the context affects both, the treatment assignment and response. In a multiple treatment scenario, we propose the neural network based **MultiMBNN**, where we overcome confounding by employing generalized propensity score based matching, and learning balanced representations. We benchmark the performance on synthetic and real-world datasets using PEHE, and mean absolute percentage error over ATE as metrics. **MultiMBNN** outperforms the state-of-the-art algorithms for CI such as **TARNet** and Perfect Match (**PM**).

1 Introduction

The primary goal in causal inference (CI) is to uncover the cause-effect relationship between entities, often construed as a problem of quantifying the effect of treatments on individuals. Randomised control trials, the most popular choice for obtaining causal relationships, are expensive and entail several logistical and ethical constraints. Causality is a crucial paradigm in several domains where observational data is available, such as healthcare [1], socioeconomic studies [2], advertising [3], etc. An impediment to CI in observational studies is the presence of *confounding*, where assignment and the response to the treatment depends on context covariates, resulting in selection bias.

In order to abate the effect of confounding, the discrepancy between the distribution of individuals receiving different treatments is minimized, in order to emulate a randomized trial. In the binary treatment case, confounding is addressed using statistical approaches such as sub-classification, weighting, imputations, and propensity score (PS) matching for unbiased per-individual causal estimates. Deep neural networks (DNN) based techniques propose counterfactual distribution modelling [4], and learning balancing representation to minimize selection bias [5, 6, 7]. In literature, the multiple treatments scenario is interpreted as different dosage levels of a single treatment, or being one of several treatments. For the latter case, matching and sub-classification techniques have been proposed [1], and in particular, generalized propensity score (GPS) based matching was proposed [8] as an accurate metric in the multiple treatment scenario [1]. DNN-based techniques for CI include counterfactual distribution learning [9], Gaussian process based modeling[10], and PS matching based Perfect match [11]. Here, we propose a novel DNN for counterfactual inference, which

overcomes confounding by leveraging on both, GPS based matching and learning balancing representations.

Contributions: We propose a novel framework for counterfactual inference in the presence of confounding due to multiple treatments. We assume strong ignorability and no hidden confounding [5, 6, 7]. We optimize DNN models via mini-batch stochastic gradient descent (SGD) to predict both, the factual and counterfactual response of a given individual to any one of the treatments. We propose the GPS based matching, along with learning balanced representations to address confounding. We estimate the GPS by training a predictive model and we use this GPS to *match* every sample within a minibatch with its nearest neighbours. In order to learn the balanced representation, we propose a loss function using the pairwise minimum mean discrepancy (MMD) metric. The novelty of this work is two-fold. First, we show that GPS based matching leads to more accurate counterfactual inference as compared to [11]. Next, we generalize the balancing representation based loss function [5, 6] to multiple treatment scenario. On synthetic and real-world datasets, we demonstrate that by combining matching and the generalized loss function, we outperform the state-of-the-art CI techniques such as perfect match (PM) [11] and TARNet [6]. We use precision in estimation of heterogeneous effect (*PEHE*) and mean absolute percentage error (*MAPE*) over the average treatment effect (*ATE*) as metrics.

2 MultiMBNN Causal Inference Model

In this section, we describe the preliminaries of the CI framework in the multiple treatment scenario, followed by details of the proposed MultiMBNN framework.

2.1 Causal Inference Preliminaries

We consider observational training data D_{CI} , comprising of N samples, where each sample is given by $\{x_i, \mathbf{t}, \mathbf{y}_i\}$. Each individual (also called *context*) i is represented using covariates given by x_i for $1 \leq i \leq N$. An individual is subjected to one of the K treatments given by $\mathbf{t} = [t_1, \dots, t_K]$, where each entry of \mathbf{t} is binary, i.e., $t_k \in \{0, 1\}$. Here, $t_k = 1$ implies that the k -th treatment is given. We assume that only one treatment is provided to an individual i at any given point in time, and hence, \mathbf{t} is an one-hot vector. Accordingly, the response vector for the i -th individual is given by $\mathbf{y}_i \in \mathbb{R}^{K \times 1}$, i.e., the outcome is a continuous random vector with K entries denoted by y_{ik} , the response of the i -th individual to the k -th treatment. We define counterfactual as the $K - 1$ alternate treatments which are unobserved for an individual.

Our goal is to train a DNN model to overcome confounding and perform counterfactual regression, i.e., to predict the response given any context and treatment. We address the issue of confounding using both, matching and learning balanced representations. In the sequel, we describe the matching method used and the loss function that caters to a multiple treatment scenario.

2.2 Generalized Propensity Score Matching

Propensity score (PS) based matching is a well-known technique used to induce the effect of randomized experiment by obtaining similar covariate distributions across treated populations [1]. Here we employ matching based on Generalised Propensity Score (GPS), which is a more relevant score for the multiple treatment scenario. PS is the conditional probability of a given individual x_i receiving a treatment t_k , i.e., $p(t_k|x_i)$. Accordingly, the GPS vector is defined as $p(\mathbf{t}|x_i) = [p(t_1|x_i), p(t_2|x_i), \dots, p(t_K|x_i)]$. In practice, we do not have access to the GPS vector, and hence, we estimate it by training a predictive model. In this work, we train an SVM or random forest. We use this tuned model to predict PS, $p(\mathbf{t}|x_i)$, on the training data used for causal inference. In order to avoid overfitting, we first obtain l nearest neighbors, and pick one out of these l samples at random for each counterfactual treatment of x_i . We employ the GPS vector $p(\mathbf{t}|x_i)$ for batch augmentation in every minibatch of SGD. For every sample within a minibatch, $K - 1$ closest neighbour samples are obtained. For instance, consider a sample x_i and its factual treatment t_i . We propose the GPS-based matching strategy M_{GPS} , which selects a neighbour x_j with observed treatment t_j such that $t_j \neq t_i$ and $d_{GPS}(i, j)$ is minimum. Here, $d_{GPS_{i,j}}$ is defined as $d_{GPS}(i, j) = \sum_{k=1}^K |p(t_k|x_i) - p(t_k|x_j)|$. On the other hand, the PS matching strategy M_{PS} [11] selects a neighbour x_j with observed treatment t_j such that $t_j \neq t_i$ and $d_{PS}(i, j)$ is minimum where $d_{PS}(i, j) = |p(t_j|x_i) - p(t_j|x_j)|$. Albeit its popularity, PS based matching has been described as inadequate and it sometimes leads to imbalance in parametric models due to model dependence [12]. Hence, we learn balancing representations to achieve better performance.

2.3 Learning Balancing Representations

In addition to overcoming the imbalance using matching, we also propose learning balanced representation using DNN. In [5], the authors perform counterfactual inference by generalizing the factual to counterfactual distribution, for the binary treatment scenario. We extend this framework from binary to a multiple treatment scenario, and modify the loss function as follows:

$$\mathcal{L}(\alpha, \gamma) = \frac{1}{N} \sum_{i=1}^N \sum_{k=1}^K |(h(\Phi(x_i), t_k) - y_{ik})| + \alpha \sum_{m=1}^K \sum_{q=1}^{m-1} \text{disc}(\hat{p}_{\Phi}^m, \hat{p}_{\Phi}^q) + \gamma \mathcal{R}(h), \quad (1)$$

where $\alpha, \gamma > 0$ are hyperparameters that control the strength of the imbalance penalties, $\mathcal{R}(h)$ is a model complexity term, $\hat{p}_{(\cdot)}^m$ and $\hat{p}_{(\cdot)}^q$ represent the distribution w.r.t. the m -th treatment and the q -th treatment, respectively, and $\text{disc}(\cdot, \cdot)$ is the minimum mean discrepancy measure as defined in [6]. We learn the balancing representation $\Phi(\cdot)$ and the hypothesis $h(\cdot)$ jointly by training a deep neural network using a loss function that incorporates the factual and the imbalance error as depicted in Fig. 1. In (1), the first term on the right hand side represents the factual loss. The second term computes the pairwise minimum mean discrepancy

between factual distributions of different treatments. The loss function in (1) is a generalisation of [6] to the multiple treatment scenario, it reduces to the one proposed in [6] for $K = 2$.

Algorithm 1 MultiMBNN algorithm

- 1: **procedure** MULTIMBNN(D)
 - 2: Split dataset D into D_{CI} for CI, and D_{PS} to compute GPS.
 - 3: Divide D_{CI} into train ($D_{CI,t}$), validation and test sets.
 - 4: Obtain GPS $p(\mathbf{t}|x_i), \forall i$ in $D_{CI,t}$, as described in Sec. 2.2.
 - 5: Divide $D_{CI,t}$ into batches with each batch being M_B
 - 6: **for** E epochs and $M_B \in D_{CI,t}$ **do**
 - 7: $\tilde{M}_B \leftarrow$ augment M_B using M_{GPS} , as described in Sec. 2.2
 - 8: Update $\Phi(\cdot), h(\cdot)$ using input \tilde{M}_B by minimizing Eq. 1
 - return** $\Phi(\cdot), h(\cdot)$
-

2.4 Proposed Approach: MultiMBNN

We propose the MultiMBNN algorithm as described in Algorithm 1. As discussed in the previous subsections, we perform batch augmentation based on GPS, and train a DNN to learn the balancing representation, Φ and the hypothesis layers h_1, \dots, h_K , one for each treatment, using the augmented minibatches, via SGD based training. The proposed neural network architecture is depicted in Fig. 1.

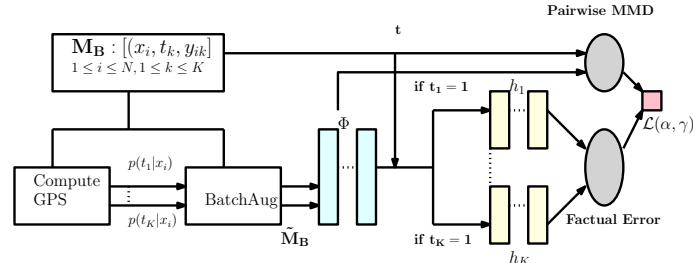


Fig. 1: MultiMBNN: Proposed Neural network architecture

3 Experiments and Results

We illustrate the performance of MultiMBNN algorithm on synthetic (named as Syn) [13], semi-synthetic NEWS [11] and cancer genome TCGA¹ datasets. We obtain $N = 15000$ samples with 10 covariates for the synthetic dataset with. Here, κ accounts for the treatment assignment bias. We employ the DGP in [11] and 5772 bag-of-words context covariates to generate $N = 10000$ samples of the NEWS dataset. In the case of synthetic and NEWS datasets, we generate data for $K = 4, 6, 8$, referred to as 'name of the dataset', followed by K . The TCGA dataset consisting of 10000 samples with 20547 covariates is obtained using the

¹https://github.com/d909b/perfect_match

DGP in [11] with $K = 4$ ('TCGA4'). We use $PEHE$ (denoted as ϵ_P) as defined in [11], and MAPE over $ATE = \frac{1}{N} \sum_{i=1}^N |y_{ik} - \frac{1}{K-1} \sum_{j=1, t_j \neq t_k}^{K-1} y_{ij}|$ [7]. We baseline the proposed algorithm using TARNet [6], MultiBNN which learns balanced representations as described in Sec. 2.3, PM is as described in [11], and MultiMBNN (M_{PS}), which uses M_{PS} (not M_{GPS}) along with balanced representation. We demonstrate the performance of MultiMBNN algorithm using several experimental settings. First, we illustrate the effect of treatment assignment bias using the parameter κ . As illustrated in Fig. 2, for Syn4 MultiMBNN performs the best for κ_3 since imbalance amongst treatments leads to one of the four treatments to be suppressed resulting in a large counterfactual error and hence, an elbow point. For NEWS4, κ_5 has the least counterfactual error since imbalance amongst treatment groups is minimum, leading to near uniform distribution of population samples in all groups.




Fig. 2: Comparison of CI frameworks: Counterfactual error across epochs (Syn4, κ_1) on the left, and $\sqrt{\epsilon_P}$ vs. κ on the right (*top right*: Syn4 with $\kappa_1 < \kappa_2 < \kappa_3 < \kappa_4$, *bottom right*: NEWS4 with $\kappa_5 < \kappa_6 < \kappa_7$).




Fig. 3: $\sqrt{\epsilon_P}$ and MAPE for Syn(κ_1) and NEWS(κ_6) datasets for varying K .

In Fig. 3, we illustrate the performance of the proposed algorithms and the baselines with varying K for a fixed κ . We see that the MultiMBNN algorithm outperforms the baselines by large margins. Further, we simulate MultiMBNN with different initial seed-points maintaining K and κ fixed, and report the mean and standard deviation in $\sqrt{\epsilon_P}$ and MAPE for all baselines in Table 1. We infer that MultiMBNN which incorporates both matching and DNN based balancing, fairs considerably well over all the baselines.

Metrics, Dataset	TARNet	MultiBNN	PM	MultiBNN(M_{PS})	MultiBNN
$\sqrt{\hat{\epsilon}_P}$, Syn4	10.21 \pm 0.56	9.34 \pm 0.61	8.21 \pm 0.35	7.98 \pm 0.23	7.86 \pm 0.37
MAPE, Syn4	0.07 \pm 0.02	0.08 \pm 0.03	0.06 \pm 0.02	0.04 \pm 0.01	0.02 \pm 0.02
$\sqrt{\hat{\epsilon}_P}$, NEWS4	9.70 \pm 1.15	9.37 \pm 0.90	9.16 \pm 0.80	8.99 \pm 0.94	8.96 \pm 0.92
MAPE, NEWS4	0.81 \pm 0.11	0.81 \pm 0.10	0.81 \pm 0.12	0.80 \pm 0.13	0.82 \pm 0.12
$\sqrt{\hat{\epsilon}_P}$, TCGA4	29.45 \pm 3.48	26.15 \pm 3.29	23.57 \pm 1.10	23.18 \pm 1.39	21.47 \pm 0.96
MAPE, TCGA4	0.93 \pm 0.14	0.84 \pm 0.10	0.92 \pm 0.07	0.78 \pm 0.03	0.80 \pm 0.08

Table 1: $\sqrt{\hat{\epsilon}_P}$ and MAPE for multiple runs of Syn4, NEWS4, TCGA4 (fixed κ).

4 Conclusions

In this work, we propose MultiBNN algorithm which addresses the inadequacies of the matching framework by learning the balanced representations in multiple treatment causal inference scenario. We demonstrate that MultiBNN outperforms the state-of-the-art techniques for multiple treatments and also for datasets with thousands of potential co-variate confounders. In future, we shall extend this algorithm for handling sparsity in the presence of large number of treatments.

References

- [1] E. A. Stuart. Matching methods for causal inference: A review and a look forward. *Statist. Sci.*, 25(1):1, 2010.
- [2] S. Athey. Machine learning and causal inference for policy evaluation. In *ACM SIGKDD*, pages 5–6. ACM, 2015.
- [3] L. Bottou, J. Peters, J. Quiñonero-Candela, D. Charles, D. M. Chickering, E. Portugaly, D. Ray, P. Simard, and E. Snelson. Counterfactual reasoning and learning systems: The example of computational advertising. *JMLR*, 14(1):3207–3260, 2013.
- [4] C. Louizos, U. Shalit, J. Mooij, D. Sontag, R. Zemel, and M. Welling. Causal effect inference with deep latent-variable models. In *NeurIPS*, pages 6446–6456, 2017.
- [5] F. Johansson, U. Shalit, and D. Sontag. Learning representations for counterfactual inference. In *ICML*, pages 3020–3029, 2016.
- [6] U. Shalit, F. D. Johansson, and D. Sontag. Estimating individual treatment effect: generalization bounds and algorithms. In *ICML*, pages 3076–3085, 2017.
- [7] A. Sharma, G. Gupta, R. Prasad, A. Chatterjee, L. Vig, and G. Shroff. MetaCI: Meta-Learning for Causal Inference in a Heterogeneous Population. *NeurIPS CausalML workshop*, 2019.
- [8] K. Imai and D. A. Van Dyk. Causal inference with general treatment regimes: Generalizing the propensity score. *JASA*, 99(467):854–866, 2004.
- [9] J. Yoon, J. Jordon, and M. van der Schaar. GANITE: Estimation of Individualized Treatment Effects using Generative Adversarial Nets. *ICLR*, 2018.
- [10] A. M. Alaa and M. van der Schaar. Bayesian inference of individualized treatment effects using multi-task gaussian processes. *NeurIPS*, pages 3424–3432, 2017.
- [11] P. Schwab, L. Linhardt, and W. Karlen. Perfect match: A simple method for learning representations for counterfactual inference with neural networks. *arXiv:1810.00656*, 2018.
- [12] G. King and R. Nielsen. Why propensity scores should not be used for matching. *Political Analysis*, 27(4):435–454, 2019.
- [13] W. Sun, P. Wang, D. Yin, J. Yang, and Y. Chang. Causal inference via sparse additive models with application to online advertising. In *AAAI*, 2015.