

Continual Causal Effect Estimation: Challenges and Opportunities

Zhixuan Chu

Ant Group, Hangzhou, China

CHUZHIXUAN.CZX@ALIBABA-INC.COM

Sheng Li

University of Virginia, Charlottesville, USA

SHENGLI@VIRGINIA.EDU

1. Introduction

A further understanding of cause and effect within observational data is critical across many domains, such as economics, health care, public policy, web mining, online advertising, and marketing campaigns. Although significant advances have been made to overcome the challenges in causal effect estimation with observational data, such as missing counterfactual outcomes and selection bias between treatment and control groups, the existing methods mainly focus on source-specific and stationary observational data. In particular, such learning strategies assume that all observational data are already available during the training phase and from only one source.

Along with the fast-growing segments of industrial applications, this assumption is unsubstantial in practice. Taking Alipay as an example, which is one of the world’s largest mobile payment platforms and offers financial services to billion-scale users, a tremendous amount of data containing much privacy-related information is produced daily and collected from different sources. In the following, we further elaborate on this problem by two points. The first one is based on the characteristics of observational data, which are incrementally available from non-stationary data distributions. For instance, electronic financial records for one marketing campaign are growing every day, and they may be collected from different cities or even other countries. This characteristic implies that one cannot have access to all observational data at one time point and from one single source. The second reason is based on the realistic consideration of accessibility. For example, when new observational data are available, one may want to refine the previously trained model using both the new data and original data. However, it is likely that the original training data are no longer accessible due to a variety of reasons, e.g., legacy data may be unrecorded, proprietary, the sensitivity of financial data, too large to store, or subject to privacy constraints of personal information (Zhang et al., 2020). This practical concern of accessibility is ubiquitous in various academic and industrial applications. That’s what it boiled down to in the era of big data, we face new challenges in causal inference with observational data. We first presented the continual causal effect estimation problem in (Chu et al., 2020a), in which we discussed three desired properties of continual causal inference frameworks, i.e., the **extensibility** for incrementally available observational data, the **adaptability** for various data sources in new domains, and the **accessibility** for an enormous amount of data.

In this position paper, we formally define the problem of continual treatment effect estimation, describe its research challenges, and then present possible solutions to this problem. Moreover, we will discuss future research directions on this topic.

2. Related Work

Instead of randomized controlled trials, observational data is obtained by the researcher simply observing the subjects without any interference. It means that the researchers have no control over the treatment assignments, and they just observe the subjects and record data based on observations (Yao et al., 2021; Chu et al., 2023a). Therefore, from the observational data, directly estimating the treatment effect is challenging due to the missing counterfactual outcomes and the existence of confounders. Recently, powerful machine learning methods such as tree-based methods (Athey and Imbens, 2016; Wager and Athey, 2018), representation learning (Li and Fu, 2017; Shalit et al., 2017; Yao et al., 2018; Chu et al., 2022), meta-learning (Künzel et al., 2019; Nie and Wager, 2021), and generative models (Louizos et al., 2017; Yoon et al., 2018) have achieved prominent progress in treatment effect estimation.

In addition, the combination of causal inference and other research fields also exhibits complementary strengths, such as computer vision (Tang et al., 2020; Liu et al., 2022a), graph learning (Ma et al., 2022; Chu et al., 2021), natural language processing (Feder et al., 2022; Liu et al., 2022b), and so on. The involved causal analysis helps improve the model’s capability of discovering and resolving the underlying system beyond the statistical relationships learned from observational data.

3. Problem Definition

Suppose that the observational data contain n units collected from d different domains, and $D_d = \{(x, y, t) | x \in X, y \in Y, t \in T\}$ denotes the dataset collected from the d -th domain, which contains n_d units. Let X denote all observed variables, Y denote the outcomes in the observational data, and T be a binary variable. Let $D_{1:d} = \{D_1, D_2, \dots, D_d\}$ be the combination of d datasets, separately collected from d different domains. For d datasets $\{D_1, D_2, \dots, D_d\}$, they have the commonly observed variables, but due to the fact that they are collected from different domains, they usually have different distributions with respect to X , Y , and T in each dataset. Each unit in the observational data received one of the two or multiple treatments. Let t_i denote the treatment assignment for unit i ; $i = 1, \dots, n$. For binary treatments, $t_i = 1$ is for the treatment group and $t_i = 0$ for the control group. The outcome for unit i is denoted by y_t^i when treatment t is applied to unit i . For observational data, only one of the potential outcomes is observed. The observed outcome is called the factual outcome, and the remaining unobserved potential outcomes are called counterfactual outcomes.

The potential outcome framework has been widely used for estimating treatment effects (Rubin, 1974; Splawa-Neyman et al., 1990). The individual treatment effect (ITE) for unit i is the difference between the potential treated and control outcomes and is defined as:

$$\text{ITE}_i = y_1^i - y_0^i. \quad (1)$$

The average treatment effect (ATE) is the difference between the mean potential treated and control outcomes, which is defined as:

$$\text{ATE} = \frac{1}{n} \sum_{i=1}^n (y_1^i - y_0^i). \quad (2)$$

The success of the potential outcome framework is based on the following assumptions (Imbens and Rubin, 2015), which ensure that the treatment effect can be identified.

Assumption 1 *Stable Unit Treatment Value Assumption (SUTVA):* The potential outcomes for any unit do not vary with the treatments assigned to other units, and, for each unit, there are no different forms or versions of each treatment level, which lead to different potential outcomes.

Assumption 2 *Consistency:* The potential outcome of treatment t is equal to the observed outcome if the actual treatment received is t .

Assumption 3 *Positivity:* For any value of x , treatment assignment is not deterministic, i.e., $P(T = t | X = x) > 0$, for all t and x .

Assumption 4 *Ignorability:* Given covariates, treatment assignment is independent of the potential outcomes, i.e., $(y_1, y_0) \perp\!\!\!\perp t | x$.

The goal of *continual treatment effect estimation* is to estimate the causal effect of treatments for all available data, including new data D_d and the previous data $D_{1:(d-1)}$, without having access to previous data $D_{1:(d-1)}$.

4. Research Challenges

Existing causal effect inference methods, however, are unable to deal with the aforementioned new challenges in continual treatment effect estimation, i.e., extensibility, adaptability, and accessibility. Although it is possible to adapt existing treatment effect estimation methods to cater to these issues, these modified methods still have inevitable defects. Three straightforward adaptation strategies are described as follows:

1. If we directly apply the model previously trained based on original data to new observational data, the performance on new tasks will be very poor due to the domain shift issues among different data sources;
2. Suppose we utilize newly available data to re-train the previously learned model for adapting changes in the data distribution. In that case, old knowledge will be completely or partially overwritten by the new one, which can result in severe performance degradation on old tasks. This is the well-known *catastrophic forgetting* problem (McCloskey and Cohen, 1989; French, 1999);
3. To overcome the catastrophic forgetting problem, we may rely on the storage of old data and combine the old and new data together, and then re-train the model from scratch. However, this strategy is memory inefficient and time-consuming, and it brings practical concerns such as copyright or privacy issues when storing data for a long time (Samet et al., 2013).

Any of these three strategies, in combination with the existing causal effect inference methods, is deficient.

5. Potential Solution

To address the continual treatment effect estimation problem, we propose a **Continual Causal Effect Representation Learning** framework (CERL) for estimating causal effect with incrementally available observational data. Instead of having access to all previous observational data, we only store a limited subset of feature representations learned from previous data. Combining selective and balanced representation learning, feature representation distillation, and feature transformation, our framework preserves the knowledge learned from previous data and updates the knowledge by leveraging new data so that it can achieve the continual causal effect estimation for incrementally new data without compromising the estimation capability for previous data. In the following, we will briefly describe the design of our CERL framework. More technical details of CERL are presented in (Chu et al., 2023b).

Framework Overview. To deal with the incrementally available observational data, the framework of CERL is mainly composed of two components: (1) the baseline causal effect learning model is only for the first available observational data, and thus we don't need to consider the domain shift issue among different data sources. This component is equivalent to the traditional causal effect estimation problem; (2) the continual causal effect learning model is for the sequentially available observational data, where we need to handle more complex issues, such as knowledge transfer, catastrophic forgetting, global representation balance, and memory constraint.

Baseline Causal Effect Learning Model. We first train the baseline causal effect learning model for the initial observational dataset and then bring in subsequent datasets. The task on the initial dataset can be converted to a traditional causal effect estimation problem. Owing to the success of deep learning for counterfactual inference, we propose to learn the selective and balanced feature representations (Shalit et al., 2017; Chu et al., 2020b) for units in treatment and control groups and then infer the potential outcomes based on learned representation space.

Sustainability of Model Learning. To avoid catastrophic forgetting when learning new data, we propose to preserve a subset of lower-dimensional feature representations rather than all original covariates. We can also adjust the number of preserved feature representations according to the memory constraint.

Continual Causal Effect Learning. We have stored memory and the baseline model. To continually estimate the causal effect for incrementally available observational data, we incorporate feature representation distillation and feature representation transformation to estimate the causal effect for all seen data based on a balanced global feature representation space.

6. Research Opportunities

Although significant advances have been made to overcome the challenges in causal effect estimation, real-world applications based on observational data are always very complicated.

Unlike source-specific and stationary observational data, most real-world data are incrementally available and from non-stationary data distributions. Significantly, we also face the realistic consideration of accessibility. Our work (Chu et al., 2020a) might be the first attempt to investigate the continual causal inference problem, and we proposed the corresponding evaluation criteria. However, constructing the comprehensive analytical tools and the theoretical framework derived from this brand-new problem requires non-trivial efforts. Specifically, there are several potential directions for continual causal inference:

- In addition to the distribution shift of the covariates among different domains, there are other potential technical issues for continual effect estimation: for example, perhaps we do not initially observe all the necessary confounding variables and may get access to increasingly more confounders.
- Compared with homogeneous treatment effects (e.g., the magnitude and direction of the treatment effect are the same for all patients, regardless of any other patient characteristics), heterogeneous causal effects could differ for different individuals. This could be another important aspect to consider for the continual treatment effect estimation model.
- The basic assumptions for traditional causal effect estimation may not be completely applicable. New assumptions may be supplemented, or previous assumptions need to be relaxed.
- There exists a natural connection with continual domain adaptation among different times or domains (“continual” causal inference) and between treatment and control groups (continual “causal inference”).
- Compared to traditional causal effect estimation tasks based on relatively small datasets, the continual causal inference method will embrace high-performance computing or cloud computing due to its ambitious objective.
- With the increasing public concern over privacy leakage in data, federated learning, which collaboratively trains the machine learning model without directly sharing the raw data among the data holders, may become a potential solution for continual causal inference.

References

- Susan Athey and Guido Imbens. Recursive partitioning for heterogeneous causal effects. *Proceedings of the National Academy of Sciences*, 113(27):7353–7360, 2016.
- Zhixuan Chu, Stephen Rathbun, and Sheng Li. Continual lifelong causal effect inference with real world evidence. 2020a.
- Zhixuan Chu, Stephen L Rathbun, and Sheng Li. Matching in selective and balanced representation space for treatment effects estimation. In *Proceedings of the 29th ACM International Conference on Information & Knowledge Management*, pages 205–214, 2020b.

- Zhixuan Chu, Stephen L Rathbun, and Sheng Li. Graph infomax adversarial learning for treatment effect estimation with networked observational data. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*, pages 176–184, 2021.
- Zhixuan Chu, Stephen L Rathbun, and Sheng Li. Learning infomax and domain-independent representations for causal effect inference with real-world data. In *Proceedings of the 2022 SIAM International Conference on Data Mining (SDM)*, pages 433–441. SIAM, 2022.
- Zhixuan Chu, Jianmin Huang, Ruopeng Li, Wei Chu, and Sheng Li. Causal effect estimation: Recent advances, challenges, and opportunities. *arXiv preprint arXiv:2302.00848*, 2023a.
- Zhixuan Chu, Ruopeng Li, Stephen Rathbun, and Sheng Li. Continual causal inference with incremental observational data. *arXiv preprint arXiv:2303.01775*, 2023b.
- Amir Feder, Katherine A Keith, Emaad Manzoor, Reid Pryzant, Dhanya Sridhar, Zach Wood-Doughty, Jacob Eisenstein, Justin Grimmer, Roi Reichart, Margaret E Roberts, et al. Causal inference in natural language processing: Estimation, prediction, interpretation and beyond. *Transactions of the Association for Computational Linguistics*, 10: 1138–1158, 2022.
- Robert M French. Catastrophic forgetting in connectionist networks. *Trends in cognitive sciences*, 3(4):128–135, 1999.
- Guido W Imbens and Donald B Rubin. *Causal inference in statistics, social, and biomedical sciences*. Cambridge University Press, 2015.
- 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.
- Sheng Li and Yun Fu. Matching on balanced nonlinear representations for treatment effects estimation. *Advances in Neural Information Processing Systems*, 30, 2017.
- Bing Liu, Dong Wang, Xu Yang, Yong Zhou, Rui Yao, Zhiwen Shao, and Jiaqi Zhao. Show, deconfound and tell: Image captioning with causal inference. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18041–18050, 2022a.
- Jiayi Liu, Wei Wei, Zhixuan Chu, Xing Gao, Ji Zhang, Tan Yan, and Yulin Kang. Incorporating causal analysis into diversified and logical response generation. In *Proceedings of the 29th International Conference on Computational Linguistics*. International Committee on Computational Linguistics, October 2022b.
- Christos Louizos, Uri Shalit, Joris M Mooij, David Sontag, Richard Zemel, and Max Welling. Causal effect inference with deep latent-variable models. In *Advances in Neural Information Processing Systems*, pages 6446–6456, 2017.

- Jing Ma, Mengting Wan, Longqi Yang, Jundong Li, Brent Hecht, and Jaime Teevan. Learning causal effects on hypergraphs. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pages 1202–1212, 2022.
- Michael McCloskey and Neal J Cohen. Catastrophic interference in connectionist networks: The sequential learning problem. In *Psychology of learning and motivation*, volume 24, pages 109–165. Elsevier, 1989.
- Xinkun Nie and Stefan Wager. Quasi-oracle estimation of heterogeneous treatment effects. *Biometrika*, 108(2):299–319, 2021.
- Donald B Rubin. Estimating causal effects of treatments in randomized and nonrandomized studies. *Journal of educational Psychology*, 66(5):688, 1974.
- Saeed Samet, Ali Miri, and Eric Granger. Incremental learning of privacy-preserving bayesian networks. *Applied Soft Computing*, 13(8):3657–3667, 2013.
- Uri Shalit, Fredrik D Johansson, and David Sontag. Estimating individual treatment effect: generalization bounds and algorithms. In *International Conference on Machine Learning*, pages 3076–3085. PMLR, 2017.
- Jerzy Splawa-Neyman, Dorota M Dabrowska, and TP Speed. On the application of probability theory to agricultural experiments. essay on principles. section 9. *Statistical Science*, pages 465–472, 1990.
- Kaihua Tang, Yulei Niu, Jianqiang Huang, Jiaxin Shi, and Hanwang Zhang. Unbiased scene graph generation from biased training. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 3716–3725, 2020.
- Stefan Wager and Susan Athey. Estimation and inference of heterogeneous treatment effects using random forests. *Journal of the American Statistical Association*, 113(523):1228–1242, 2018.
- Liuyi Yao, Sheng Li, Yaliang Li, Mengdi Huai, Jing Gao, and Aidong Zhang. Representation learning for treatment effect estimation from observational data. *Advances in Neural Information Processing Systems*, 31, 2018.
- Liuyi Yao, Zhixuan Chu, Sheng Li, Yaliang Li, Jing Gao, and Aidong Zhang. A survey on causal inference. *ACM Transactions on Knowledge Discovery from Data (TKDD)*, 15(5): 1–46, 2021.
- Jinsung Yoon, James Jordon, and Mihaela van der Schaar. Ganite: Estimation of individualized treatment effects using generative adversarial nets. In *International Conference on Learning Representations*, 2018.
- Junting Zhang, Jie Zhang, Shalini Ghosh, Dawei Li, Serafettin Tasci, Larry Heck, Heming Zhang, and C-C Jay Kuo. Class-incremental learning via deep model consolidation. In *The IEEE Winter Conference on Applications of Computer Vision*, pages 1131–1140, 2020.