Contribution information sheet for
Partial counterfactual identification and uplift modeling: theoretical results and real-world assessment

Théo Verhelst, Denis Mercier,
Jeevan Shrestha, Gianluca Bontempi

March 2022

Abstract

Counterfactuals are central in causal human reasoning and the scientific discovery process. The uplift, also called conditional average treatment effect, measures the causal effect of some action, or treatment, on the outcome of an individual. This paper discusses how it is possible to derive bounds on the probability of counterfactual statements based on uplift terms. First, we derive some original bounds on the probability of counterfactuals and we show that tightness of such bounds depends on the information of the feature set on the uplift term. Then, we propose a point estimator based on the assumption of conditional independence between the counterfactual outcomes. The quality of the bounds and the point estimators are assessed on synthetic data and a large real-world customer data set provided by a telecom company, showing significant improvement over the state of the art.

Contributions

The contributions of this paper are as follow:

- A set of original bounds on the probability of counterfactuals, expressed in terms of the uplift quantity.
- A formal derivation of the relationship between our original bounds and the state-of-the-art Frechet bounds derived by Tian & Pearl (2000).
- A point estimator of the counterfactual probabilities based on the conditional independence assumption.
- A hierarchical Bayesian model for simulating counterfactual settings and assessing the accuracy of the sample version of the derived bounds.
- A real-world assessment of the proposed bounds with a large data set of customer churn campaigns and a discussion of the potential benefits.

Simulated examples indicate that the proposed bounds typically provide a significant improvement over the state of the art, and that the point estimator provides a good approximation of the true counterfactual probability even when the underlying assumption is not respected.

Related work

The *probability of necessity and sufficiency* (PNS) as presented by Pearl (2009, p. 286) is one of the four counterfactual probabilities that we consider in this paper. Seminal works on partial counterfactual identification include (Balke & Pearl, 1994) and (Tian & Pearl, 2000). The PNS conditioned on a set of covariates x is called x-specific PNS in (Li & Pearl, 2019). The main focus of Li & Pearl (2019) is the estimation of the benefit generated by a customer retention campaign when the different types of customers have different values. In (Li & Pearl, 2021), the authors further refine the bounds on the campaign benefit based on causal assumptions derived from causal diagrams.

Mueller et al. (2021) derived tighter bounds on the PNS for a variety of causal diagrams, such as with sufficient covariates or with a mediator variable. In particular, Theorem 5 in (Mueller et al., 2021) is formally very close to the bounds we develop in this paper, although they consider a set of discrete covariates, whereas we use uplift modeling which allows for arbitrary high-dimensional covariate sets. Zhang et al. (2022) express the problem of bounding the probability of counterfactuals into polynomial programming, providing tight bounds for any causal graph and combination of experimental and observational data.

Our approach in this paper differs from Mueller et al. (2021) and Zhang et al. (2022) in that we make very few causal assumptions (only that the treatment is randomized), but we suggest uplift modeling as a powerful way to estimate conditional probabilities, and we analyze the impact of mutual information between the conditioning set and the potential outcomes.

Past submissions

This paper is the revised version of our original submission to the MLJ special issue on the Foundations of Data Science[1]. The manuscript has been significantly improved following the reviewers' comments. In particular, we developed further our theoretical contributions on bounding and approximating the probability of counterfactuals (Sections 4 and 5). We added two theorems to set these results on firm theoretical grounds. We separated the general results from their applications to customer churn, which required changing the name of the manuscript to "Partial counterfactual identification and uplift modeling: theoretical results and real-world assessment". Most of the text has been updated to improve clarity, and a sensitivity analysis has been added in Section 6.

References

Balke, A. & Pearl, J. (1994). Counterfactual probabilities: Computational methods, bounds and applications. In *Uncertainty Proceedings 1994* (pp. 46–54). Elsevier.

Li, A. & Pearl, J. (2019). Unit Selection Based on Counterfactual Logic. In *IJCAI* (pp. 1793–1799).: International Joint Conferences on Artificial Intelligence Organization.

[1] http://dsaa2022.dsaa.co/mlj-special-issue-on-foundations-of-data-science/

Li, A. & Pearl, J. (2021). Unit Selection with Causal Diagram. *arXiv preprint arXiv:2109.07556*.

Mueller, S., Li, A., & Pearl, J. (2021). Causes of Effects: Learning individual responses from population data. *arXiv preprint arXiv:2104.13730*.

Pearl, J. (2009). *Causality: models, reasoning, and inference*. Cambridge university press.

Tian, J. & Pearl, J. (2000). Probabilities of causation: Bounds and identification. *Annals of Mathematics and Artificial Intelligence*, 28(1), 287–313. Publisher: Springer.

Zhang, J., Tian, J., & Bareinboim, E. (2022). Partial counterfactual identification from observational and experimental data. In *International conference on machine learning* (pp. 26548–26558). tex.organization: PMLR.