Unit Selection: Case Study and Comparison with A/B Test Heuristic

Anonymous Author(s)
Affiliation
Address
email

Abstract

The unit selection problem aims to identify a set of individuals who are most likely to exhibit a desired mode of behavior, for example, selecting individuals who would respond one way if encouraged and a different way if not encouraged. Using a combination of experimental and observational data, Li and Pearl derived tight bounds on the "benefit function", which is the payoff/cost associated with selecting an individual with given characteristics. In this study, we explain the relationships between the benefit function and the A/B test heuristics by analyzing the essence of the latter one. We further show the A/B test heuristics failed in some common real-world scenarios using simulated case studies.

1 Introduction

Several areas of industry, marketing, and health science face the unit selection dilemma. For example, in customer relationship management [1, 6, 7, 16], it is useful to determine the customers who are going to leave but might reconsider if encouraged to stay. Due to the high expense of such initiatives, management is forced to limit inducement to customers who are most likely to exhibit the behavior of interest. As another example, companies are interested in identifying users who would click on an advertisement if and only if it is highlighted in online advertising [2, 10, 13, 15, 17]. The challenge in identifying these users stems from the fact that the desired response pattern is not observed directly but rather is defined counterfactually in terms of what the individual would do under hypothetical unrealized conditions. For example, when we observe that a user has clicked on a highlighted advertisement, we do not know whether they would click on that same advertisement if it were not highlighted.

The benefit function for the unit selection problem was defined by Li and Pearl [11], and it properly captures the nature of the desired behavior. Using a combination of experimental and observational data, Li and Pearl derived tight bounds of the benefit function. The only assumption is that the treatment has no effect on the population-specific characteristics. Inspired by the idea of Mueller, Li, and Pearl [14] and Dawid et al. [3] that the bounds of probabilities of causation could be narrowed using covariates information, Li and Pearl [12] narrowed the bounds of the benefit function using covariates information and their causal structure. Besides, the unit selection problem with nonbinary treatment and effect was studied by Li and Pearl in [8, 9].

In Li-Pearl's work [11], they stated that the A/B test heuristics are sometimes problematic. In this study, we focus on the relationship between Li-Pearl's benefit function and the A/B test heuristics, explaining why the A/B test heuristics are sometimes problematic. We then provide more simulated case studies to emphasize the conclusion, as well as illustrate how to apply Li-Pearl's unit selection model correctly.

Submitted to 36th Conference on Neural Information Processing Systems (NeurIPS 2022). Do not distribute.

2 Preliminaries

In this section, we review Li and Pearl's benefit function of the unit selection problem [11].

In this study we use the language of counterfactuals in structural model semantics, as given in [4, 5]. we use $Y_x = y$ to denote the counterfactual sentence "Variable Y would have the value y, had X been x". For simplicity purposes, in the rest of the paper, we use y_x to denote the event $Y_x = y$, $y_{x'}$ to denote the event $Y_{x'} = y$, y'_x to denote the event $Y_x = y'$, and $y'_{x'}$ to denote the event $Y_{x'} = y'$. we assume that experimental data will be summarized in the form of the causal effects such as $P(y_x)$ and observational data will be summarized in the form of the joint probability function such as $P(x, y)$. If not specified, the variable X stands for treatment and the variable Y stands for effect.

Individual behavior was classified into four response types: labeled complier, always-taker, never-taker, and defier. Suppose the benefit of selecting one individual in each category are $\beta, \gamma, \theta, \delta$ respectively (i.e., the benefit vector is $(\beta, \gamma, \theta, \delta)$). Li and Pearl defined the objective function of the unit selection problem as the average benefit gained per individual. Suppose x and x' are binary treatments, y and y' are binary outcomes, and c are population-specific characteristics, the objective function (i.e., benefit function) is following (If the goal is to evaluate the average benefit gained per individual for a specific population c, $argmax_c$ can be dropped.):

$$argmax_c \; \beta P(y_x, y'_{x'}|c) + \gamma P(y_x, y_{x'}|c) + \theta P(y'_x, y'_{x'}|c) + \delta P(y'_x, y_{x'}|c). \tag{1}$$

Using a combination of experimental and observational data, Li and Pearl established the most general tight bounds on this benefit function as follow (which we refer to as Li-Pearl's Theorem in the rest of the paper). The only constraint is that the population-specific characteristics are not a descendant of the treatment.

Theorem 1. *Given a causal diagram G and distribution compatible with G, let C be a set of variables that does not contain any descendant of X in G, then the benefit function $f(c) = \beta P(y_x, y'_{x'}|c) + \gamma P(y_x, y_{x'}|c) + \theta P(y'_x, y'_{x'}|c) + \delta P(y_{x'}, y'_x|c)$ is bounded as follows:*

$$\begin{aligned} W + \sigma U \le f(c) \le W + \sigma L \qquad & \text{if } \sigma < 0, \\ W + \sigma L \le f(c) \le W + \sigma U \qquad & \text{if } \sigma > 0, \end{aligned}$$

where σ, W, L, U are given by,

$$\sigma = \beta - \gamma - \theta + \delta, W = (\gamma - \delta)P(y_x|c) + \delta P(y_{x'}|c) + \theta P(y'_{x'}|c),$$

$$L = \max \left\{ \begin{array}{c} 0, \\ P(y_x|c) - P(y_{x'}|c), \\ P(y|c) - P(y_{x'}|c), \\ P(y_x|c) - P(y|c) \end{array} \right\}, U = \min \left\{ \begin{array}{c} P(y_x|c), \\ P(y'_{x'}|c), \\ P(y, x|c) + P(y', x'|c), \\ P(y_x|c) - P(y_{x'}|c) + \\ +P(y, x'|c) + P(y', x|c) \end{array} \right\}.$$

Li and Pearl also provided conditions such that the benefit function can have a point estimation.

Definition 2. *(Monotonicity) A Variable Y is said to be monotonic relative to variable X in a causal model M iff*

$$y'_x \wedge y_{x'} = \textit{false}$$

Definition 3. *(Gain Equality) The benefit of selecting a complier (β), an always-taker (γ), a never-taker(θ), and a defier (δ) is said to satisfy gain equality iff*

$$\beta + \delta = \gamma + \theta$$

Theorem 4. *Given that Y is monotonic relative to X or that $(\beta, \gamma, \theta, \delta)$ satisfies gain equality, the benefit function $f(c)$ is given by*

$$f(c) = (\beta - \theta)P(y_x|c) + (\gamma - \beta)P(y_{x'}|c) + \theta$$

3 The Essence of A/B Test Heuristic

A common solution that is explored in the literature is an A/B-test-based approach, where a controlled experiment is performed and the result is used as a criterion for selection. Specifically, individuals are randomly split into two groups called control and treatment. Individuals in the control group are served no treatment, whereas those in the treatment group are served the treatment. Then, the commonly used A/B test heuristics are $aP(y_x|c) - bP(y_{x'}|c)$ (i.e., the weighted difference between the effective rate under treatment and the effective rate under no treatment.) We then have,

$$\begin{aligned} aP(y_x|c) - bP(y_{x'}|c) &= a(P(y_x, {y'}_{x'}|c) + P(y_x, y_{x'}|c)) - b(P(y_{x'}, {y'}_x|c) + P(y_{x'}, y_x|c)) \\ &= aP(y_x, {y'}_{x'}|c) + (a-b)P(y_x, y_{x'}|c) - bP(y_{x'}, {y'}_x|c). \end{aligned}$$

The effective rate under treatment is the percentage of complier plus always-taker in the population, and the effective rate under no treatment is the percentage of always-taker plus defier in the population. The essence of A/B test heuristics is a weighted difference between (complier+always_taker) and (always_taker+defier). Therefore, the A/B test heuristics are special cases of the benefit function. The benefit function has more expression power than the A/B test heuristics. This explains why A/B test heuristics can be optimal for some cases (i.e., Gain equality [11] satisfied.) and problematic in general.

4 Case Studies

The benefit vector in Li-Pearl's model is not determined by the model but by the one who uses the model. In this section, we illustrate several common applications showing how to set the benefit vector. We categorize the applications based on the quality of A/B-test-based approaches.

4.1 Cases in which A/B-test Heuristics are Correct

4.1.1 Number of Increased Customers

Consider a mobile carrier that wants to identify customers likely to discontinue their services within the next quarter based on customer characteristics (the company management has access to user data, such as income, age, usage, and monthly payments). The carrier will then offer these customers a special renewal deal to dissuade them from discontinuing their services and to increase their service renewal rate.

Let $A = a$ denote the event that a customer receives the special deal, $A = a'$ denote the event that a customer receives no special deal, $R = r$ denote the event that a customer continues the services, $R = r'$ denote the event that a customer discontinues the services, and C (a set of variables) denote the population-specific characteristics of a customer (e.g., income, age, usage, and monthly payments).

If the manager only wants to maximize the number of increased customers due to the offer in the next quarter regardless of the total profit, then they should assign 1 to a complier because the company gains one customer due to the offer, assign 0 to an always-taker and a never-taker because the company gains no customer due to the offer, and assign -1 to a defier because the company loses one customer due to the offer.

Therefore, the benefit vector above is $(1, 0, 0, -1)$, and using Theorem 4, when the benefit vector satisfies the gain equality $(1 - 1 = 0 + 0)$, the benefit function is $f(c) = P(r_a|c) - P(r_{a'}|c)$. This is the most common A/B test heuristic in literature. From the view of the essence of A/B test heuristic, $P(r_a|c) - P(r_{a'}|c)$ is (complier+always_taker)-(always_taker+defier)=complier-defier.

4.1.2 Number of Total Customers

If the manager only wants to maximize the total number of customers in the next quarter regardless of the total profit, then they should assign 1 to a complier and an always-taker because the company has one customer in the next quarter and assign 0 to a never-taker and a defier because the company has no customer in the next quarter.

Therefore, the benefit vector above is $(1, 1, 0, 0)$, and using Theorem 4, when the benefit vector satisfies the gain equality $(1 + 0 = 1 + 0)$, the benefit function is $f(c) = P(r_a|c)$. This is another common A/B test heuristic in literature, which is the causal effect of the offer to the number of customers. From the view of the essence of A/B test heuristics, $P(r_a|c)$ is exactly complier+always_taker.

4.1.3 Immediate Profit

If the manager wants to maximize the total immediate profit due to the offer. The management estimates that the benefit of selecting a complier is $\$100$ as the profit is $\$140$ but the discount is $\$40$, the benefit of selecting an always-taker is $-\$40$ as the customer would continue the service anyway and the company loses the value of the discount, the benefit of selecting a never-taker is $\$0$ as the cost of issuing the discount is negligible, and the benefit of selecting a defier is $-\$140$ as they lose a customer due to the special offer.

Therefore, the benefit vector above is $(100, -40, 0, -140)$, using Theorem 4, when the benefit vector satisfies the gain equality $(100 - 140 = -40 + 0)$, the benefit function is $f(c) = 100P(r_a|c) - 140P(r_{a'}|c)$. This result is the same as the popular method in the industry, which is called revenue difference. The profit of a continuing customer if issued the special offer is $\$100$ and the profit of a continuing customer if no special offer is issued is $\$140$; therefore, the revenue difference is $100P(r_a|c) - 140P(r_{a'}|c)$. From the view of the essence of A/B test heuristic, $100P(r_a|c) - 140P(r_{a'}|c)$ is 100(complier+always_taker)-140(always_taker+defier)=100complier-40always_taker-140defier.

4.2 Cases in which A/B-test Heuristics are not Correct

4.2.1 Nonimmediate Profit

If the manager wants to maximize the total profit including the nonimmediate profit due to the offer. The management estimates that the benefit of selecting a complier is $\$100$ as the profit is $\$140$ but the discount is $\$40$, the benefit of selecting an always-taker is $-\$60$ as the customer would continue the service anyway (so the company loses the value of the discount and an extra cost $\$20$ because the always-taker may require additional discounts in the future), the benefit of selecting a never-taker is 0 as the cost of issuing the discount is negligible, and the benefit of selecting a defier is $-\$140$ as they lose a customer due to the special offer.

Therefore, the benefit vector above is $(100, -60, 0, -140)$, and this is the example Li-Pearl have illustrated in [11], where the simple A/B-test-based approach is NOT correct.

4.2.2 Maximize Users Satisfaction

The management of a search engine company wants to decide whether it is worth sending an advertisement to a group of users, so as to maximize overall satisfaction. The management estimates that the satisfaction of recommending an advertisement to a complier is 2 degrees, as users would gain new information that they needed, that of recommending the advertisement to an always-taker is 1 degree, as users got a shortcut to the advertisement, that of recommending the advertisement to a never-taker is -1 degrees, as users got unnecessary information, and that of recommending the advertisement to a defier is -2 degrees, as the recommendation would prevent users to get needed information (compliers are the users who would click on the advertisement if the advertisement is recommended and would not if otherwise; always-takers are the users who would click on the advertisement whether or not the advertisement is recommended; never-takers are the users who would not click on the advertisement whether or not the advertisement is recommended; defiers are the users who would click on the advertisement if the advertisement is not recommended and would not if otherwise).

Therefore, the benefit vector above is $(2, 1, -1, -2)$, and this is another example Li-Pearl have illustrated in [11], where a simple A/B-test-based approach is NOT correct because the coefficients are difficult to be determined.

4.2.3 Maximize Difference between the Number of Effective Patients and the Number of Ineffective Patients

A pharmaceutical factory invents a new medicine and wants to identify patients so as to maximize difference between the number of effective patients and the number of ineffective patients.

Table 1: Results of a simulated study on patients.

	Group1 with r	Group1 with r'	Group2 with r	Group2 with r'
$do(a)$	210	140	217	133
$do(a')$	105	245	129	221

Table 2: Results of the two objective functions based on the data from the simulated study.

	f_1	f_2	real
Group 1	0.3	-0.1	-0.2
Group 2	0.25	0.14	0.2

Therefore, they should assign 1 to a complier because the complier is the patient cured by the medicine, assign -1 to an always-taker, a never-taker, and a defier because they are all ineffective patients. The benefit vector is then $(1, -1, -1, -1)$.

Let $A = a$ denote the event that a patient receives the medicine, $A = a'$ denote the event that a patient receives no medicine, $R = r$ denote the event that a patient is cured, $R = r'$ denote the event that a patient is not cured, and C (a set of variables) denote the population-specific characteristics of a patient.

Suppose they have two groups of patients, group 1 with characteristics c_1 and group 2 with characteristics c_2. In addition, they have prior information that $P(r|c_1) = 0.3$ and $P(r|c_2) = 0.1$. They randomly select 700 patients from each group and offer the medicine to 350 customers in each group. Table 1 summarizes the results.

Let us compare the two selection strategies, each using a different objective function. The first is a simple A/B test heuristic, that is:

$$Obj_1 = argmax_c f_1(c) = argmax_c P(r|c, do(a)) - P(r|c, do(a')).$$

The second is the proposed approach, that is:

$$Obj_2 \quad = \quad argmax_c f_2(c) = argmax_c P(r_a, r'_{a'}|c) - P(r_a, r_{a'}|c) - P(r'_a, r'_{a'}|c) - P(r'_a, r_{a'}|c).$$

Then, we enter the data in Table 1 into the objective functions of groups 1 and 2. Table 2 summarizes the results (note that we use the midpoint of the bounds from Theorem 1 as the selection criterion for Obj_2). The proposed approach selected group 2; however, the first objective function selected group 1 as the desired patients.

An informer with access to the fractions of compliers, always-takers, never-takers, and defiers in both groups (as summarized in Table 3, and these numbers are never known in reality) would easily conclude that the A/B test heuristic had reached a wrong conclusion. In detail, the expected benefit of selecting a patient in group 1 is $1 \times 0.4 - 1 \times 0.2 - 1 \times 0.3 - 1 \times 0.1 = -0.2$, which means offering the medicine to group 1 would have negative difference; the expected benefit of selecting a patient in group 2 is $1 \times 0.6 - 1 \times 0.02 - 1 \times 0.03 - 1 \times 0.35 = 0.2$. Thus, the pharmaceutical factory should only offer the medicine to group 2.

5 Conclusion

We reviewed Li-Pearl's unit selection model and its benefit function. We explained the relationships between the benefit function and the A/B test heuristics by showing the essence of the latter one is a weighted difference between complier+always_taker and always_taker+defier. We further provided more simulated examples to show when the A/B test heuristics failed and how to apply Li-Pearl's model correctly.

Table 3: Percentages of four response types in each group for patients.

	Complier	Always-taker	Never-taker	Defier
Group 1	40%	20%	30%	10%
Group 2	60%	2%	3%	35%

References

[1] Alex Berson, Stephen Smith, and Kurt Thearling. *Building data mining applications for CRM*. McGraw-Hill Professional, 1999.

[2] Léon Bottou, Jonas Peters, Joaquin Quiñonero-Candela, Denis X Charles, D Max Chickering, Elon Portugaly, Dipankar Ray, Patrice Simard, and Ed Snelson. Counterfactual reasoning and learning systems: The example of computational advertising. *The Journal of Machine Learning Research*, 14(1):3207–3260, 2013.

[3] Philip Dawid, Monica Musio, and Rossella Murtas. The probability of causation. *Law, Probability and Risk*, (16):163–179, 2017.

[4] David Galles and Judea Pearl. An axiomatic characterization of causal counterfactuals. *Foundations of Science*, 3(1):151–182, 1998.

[5] Joseph Y Halpern. Axiomatizing causal reasoning. *Journal of Artificial Intelligence Research*, 12:317–337, 2000.

[6] Shin-Yuan Hung, David C Yen, and Hsiu-Yu Wang. Applying data mining to telecom churn management. *Expert Systems with Applications*, 31(3):515–524, 2006.

[7] Miguel APM Lejeune. Measuring the impact of data mining on churn management. *Internet Research*, 11(5):375–387, 2001.

[8] A. Li and J. Pearl. Probabilities of causation with non-binary treatment and effect. Technical Report R-516, Department of Computer Science, University of California, Los Angeles, CA, 2022.

[9] A. Li and J. Pearl. Unit selection with nonbinary treatment and effect. Technical Report R-517, <http://ftp.cs.ucla.edu/pub/stat_ser/r517.pdf>, Department of Computer Science, University of California, Los Angeles, CA, 2022.

[10] Ang Li, Suming J. Chen, Jingzheng Qin, and Zhen Qin. Training machine learning models with causal logic. In *Companion Proceedings of the Web Conference 2020*, pages 557–561, 2020.

[11] Ang Li and Judea Pearl. Unit selection based on counterfactual logic. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI-19*, pages 1793–1799. International Joint Conferences on Artificial Intelligence Organization, 7 2019.

[12] Ang Li and Judea Pearl. Unit selection with causal diagram. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 5765–5772, 2022.

[13] Lihong Li, Shunbao Chen, Jim Kleban, and Ankur Gupta. Counterfactual estimation and optimization of click metrics for search engines. *arXiv preprint arXiv:1403.1891*, 2014.

[14] S. Mueller, A. Li, and J. Pearl. Causes of effects: Learning individual responses from population data. Technical Report R-505, <http://ftp.cs.ucla.edu/pub/stat_ser/r505.pdf>, Department of Computer Science, University of California, Los Angeles, CA, 2021. Forthcoming, Proceedings of IJCAI-2022.

[15] Wei Sun, Pengyuan Wang, Dawei Yin, Jian Yang, and Yi Chang. Causal inference via sparse additive models with application to online advertising. In *AAAI*, pages 297–303, 2015.

[16] Chih-Fong Tsai and Yu-Hsin Lu. Customer churn prediction by hybrid neural networks. *Expert Systems with Applications*, 36(10):12547–12553, 2009.

[17] Jun Yan, Ning Liu, Gang Wang, Wen Zhang, Yun Jiang, and Zheng Chen. How much can behavioral targeting help online advertising? In *Proceedings of the 18th international conference on World Wide Web*, pages 261–270. ACM, 2009.

Checklist

1. For all authors...
   (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
   (b) Did you describe the limitations of your work? [Yes]
   (c) Did you discuss any potential negative societal impacts of your work? [N/A]
   (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
   (a) Did you state the full set of assumptions of all theoretical results? [Yes]
   (b) Did you include complete proofs of all theoretical results? [Yes] in the appendix
3. If you ran experiments...
   (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes]
   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [N/A]
   (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [Yes]
   (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [N/A]
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
   (a) If your work uses existing assets, did you cite the creators? [N/A]
   (b) Did you mention the license of the assets? [N/A]
   (c) Did you include any new assets either in the supplemental material or as a URL? [N/A]
   (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]
   (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]
5. If you used crowdsourcing or conducted research with human subjects...
   (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
   (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
   (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]