Causal AI Framework for Unit Selection in Optimizing Electric Vehicle Procurement

Chi Zhang[1], Ang Li[2], Scott Mueller[3], Rumen Iliev[1]

[1]Toyota Research Institute
[2]Department of Computer Science, Florida State University
[3]Cognitive Systems Laboratory, Department of Computer Science, University of California, Los Angeles
chi.zhang.ctr@tri.global, angli@cs.fsu.edu, scott@cs.ucla.edu, rumen.iliev@tri.global

Abstract

Electric vehicles (EVs) are generally considered more environmentally sustainable than internal combustion engine vehicles (ICEVs). Government and policy makers may want to incentivize multi-vehicle households who, if they purchase a new EV, would use their EV to replace a large portion of their ICEV mileage. Therefore, it is important to analyze how EV procurement affects annual EV mileage for different households. Given that many relevant data, especially experimental data, are often unavailable in the real world, we need causal analysis tools to answer this question. Additionally, our aim is to compare the expected EV mileage of different combinations of vehicles a household owns. Observing multiple combinations in an individual household is impossible since only one combination can exist, making causal inference challenging. In this paper, we construct a causal AI framework utilizing counterfactual reasoning methods to address this issue.

Introduction

The transportation industry contributes to more than a quarter of total greenhouse gas (GHG) emissions in the United States, and light-duty vehicles alone are responsible for more than half of these emissions (EPA 2023). There is a widespread consensus that the adoption of electrified vehicles will be a significant factor in future initiatives to achieve carbon neutrality (Jenn 2020; Burnham et al. 2021). When targeting individual choices and when the interventions have a corresponding cost associated with them it is important to take into account the possibility for heterogeneous treatment effects. The benefits from intervening on some groups, or on some individuals, might be smaller or larger than the benefits of intervening on other groups or individuals. Understanding the heterogeneity of driving patterns across individuals, households and groups is important when trying to maximize the desired outputs. A recent paper (Nunes, Woodley, and Rossetti 2022) compares the benefits from from targeting different types of households. The main difference between households the authors considered in their model was in the number of current vehicles in a household. The results indicated that the advantages of acquiring an EV could drastically vary depending on the current vehicle mix.

Nowadays many households own more than one vehicle. In particular, many people choose to purchase an EV as a complementary vehicle, not driving it much while primarily relying on their ICEV (Burlig et al. 2021). This could be due to various reasons including their personal preference towards their ICEVs, insufficient EV mileage ranges, and charging inconvenience. As a result, the carbon emission benefit is not as large as households who drive their EVs as primary vehicles. In the interest of budget, policy makers may want to target EV purchase incentives on those who, upon purchasing new EVs, would use their EVs to replace a large portion of their ICEV driving mileage. To solve this optimization problem, we need to answer the question, "what is the expected difference in EV mileage among households convinced to purchase a new EV versus not convinced?" This gives policy makers a useful criterion for prioritizing incentives.

Note that this question, at the individual household level, is counterfactual. We can never observe or test both actions, one of them cannot occur. There are significant caveats with not treating this at the individual level (Mueller and Pearl 2022). Li and Pearl detail the sometimes severely suboptimal decision making that results from a traditional analysis (Li and Pearl 2019).

In this paper, we focus on multi-vehicle households, and develop a causal AI framework to estimate the counterfactual effects of adding an additional EV to a household on the increment of their EV driving mileage.

Preliminaries

Causal Inference

A causal model is composed of a causal directed acyclic graph (DAG) $G(V, E)$ and a set of structural equations. V are nodes representing model variables and E are edges representing causal relations between two nodes. Directed edges encode the direction of causality, i.e., if a variable A is in the structural equation that determines another variable B, an edge is drawn from A to B.

In this paper, we follow the notation in (Pearl 2009), and use uppercase letters to denote variables, and lowercase letters (combined with symbols and numbers) to denote the values a variable can take on. For example, the values of a binary variable A can be denoted as a and a', and the values of a non-binary variable B can be denoted as $b^1, b^2, \ldots, b^n$. A variable C with subscript $D = d$ represents the event of C with the intervention $D = d$. This is denoted as $C_{D=d}$; where appropriate, C_d is used as a shorthand. In addition, when the time a variable is measured is not the time of interest, we will parenthesize the variable with an offset from baseline time t; i.e. $X(t-2)$ represents the variable X measured 2 time units in the past.

Causal AI Framework for EV Driving Analysis

Causal Model

We are interested in learning about what happens to a household's total miles on EV if they purchase an additional EV. We first build a graphical representation of the causal relationships using a DAG.

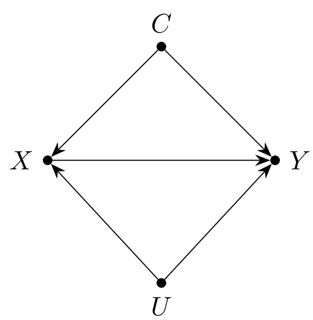

Figure 1: Causal DAG with observed confounder C and unobserved confounder U.

Since our analysis is not based on real-world data and is for illustration purposes, we limit our focus by assuming the variables are categorical and can only take on specific values. Our approach can be easily generalized for larger sets of values. For the same reason, we model one observed confounder and one unobserved confounder while this framework applies to more confounders of either type.

In this model, variable X represents the numbers and types of cars a household owns. We focus our discussion on cases where a household has 1) two EVs and one ICEV or 2) one EV and one ICEV. Cases 1 and 2 are represented by $X = x$ and $X = x'$, respectively. Variable Y represents the annual miles driven on all the EVs for a household, which lies in one of the five ranges: $y^1 = [0, 5000)$, $y^2 = [5000, 10000)$, $y^3 = [10000, 15000)$, $y^4 = [15000, 20000)$, $y^5 = [20000, \infty)$. Variables C and U are two confounders causing both X and Y. C is observed (assumed available in the data), which represents the annual total travel needs for a household using all available vehicles. $C = c$ or $C = c'$ denotes a household needs to travel more or less than 15000 miles per year, respectively. U is unobserved (assumed not available in the data), which represents whether the typical trip types of a household favors EVs. $U = u$ indicates that the household mainly drives trips that favor EVs (e.g., shorter trips or trips with easy charging options), and $U = u'$ indicates that the household mainly drives trips that favor ICEVs (e.g., longer trips or trips to cold places). Note that, although U is unobserved for single households, we might have an estimate of what percentage of the population has $U = u$ or $U = u'$. Hence, we assume that the prior $P(U)$ is given.

Problem Setting

We are interested in selecting households from a population which may benefit the environment the most from adding an additional EV. There are two aspects to this, selecting households with a larger difference in average treatment effects, or selecting households with certain "response types" (Li and Pearl 2019). For the first aspect, we want to optimize selection of households that will replace the most of their ICEV mileages. For the second aspect, we optimize selection by taking into consideration how each household responds to the treatment. For example, selecting a household whose EV usage will decrease with the addition of a new EV is strongly not preferred (and is weighted negatively). Similarly, selecting households with a larger EV usage increment is strongly preferred (and is weighted positively). One major challenge is that the unit response type is unobserved. This is known as the fundamental problem of causal inference, namely, that we can only ever observe a single outcome for an individual. We can, however, use observations from the past to gain information about the unit response type. We discuss the two aspects separately next.

Aspect 1: We are interested in assessing for a household their expected annual EV miles driven if an additional EV is added to their household fleet, where we observe they drove one EV and one ICEV that drove a certain number of EV miles for the past year. To put this into a causal expression (Pearl 2009), we have

$$P(Y_{X=x} = y^a \mid X(t-1) = x', Y(t-1) = y^p), \quad (1)$$

where $X(t-1)$ and $Y(t-1)$ are the variables X and Y at the previous timestamp, when the observational data are given. y^a and y^p are values of Y and $Y(t-1)$, respectively, with $1 \leq a, p \leq 5$. Y and $Y(t-1)$ are different variables that take on the same set of values. Here, our goal is to estimate, for a household with one EV and one ICEV that drove y^p EV miles last year, what the probability is that they would drive y^a this year if they added a new EV. We want to estimate this probability for all a and p.

Aspect 2: We are interested in selecting certain households depending on their response types. For instance, we want to know if a household would respond a certain way with one intervention, and would respond another way with another intervention. This is a counterfactual question because when we observe a household with some combination of vehicles (e.g., one EV and one ICEV), we do not simutaneously observe them with a different combination of vehicles (e.g., two EVs and one ICEV).To put this into a counterfactual expression (Pearl 2009; Li and Pearl 2022a), we have

$$P(Y_{X=x} = y^a, Y_{X=x'} = y^b | X(t-1) = x', Y(t-1) = y^p), \quad (2)$$

where $X(t-1)$ and $Y(t-1)$ are the variables X and Y at the previous timestamp where the observational data are given. y^a, y^b, y^p are values of Y (or $Y(t-1)$) with $1 \leq a, b, p \leq 5$. Y and $Y(t-1)$ are different variables but take on the same set of values. Note that this expression is the non-binary probability of necessity and sufficiency (PNS(2)) (Li and Pearl 2022a) of X on Y. Here, our goal is to estimate, for a household with one EV and one ICEV that drove y^p EV miles last year, what the probability is that they would drive y^b this year if they do not add a new EV, and would drive y^a if they do add a new EV. We want to estimate this probability for all a, b, and p.

Estimation

Aspect 1: Average Treatment Effect of Subpopulations One challenge of this problem is the unavailability of experimental data. It is prohibitively costly to conduct an experiment to provide households with EVs. Fortunately, we can use observational data to provide insights on what would be experimental results. Another challenge of causal estimation in this case is that the observational data are from the past, while we are trying to infer the behaviors for the future. To make use of the available observational data, assumptions need to be made about how past observational data predicts future states. Causal inference frameworks, when applied to real-world problems, often implicitly assume that what we observe in the past continues to apply in the future. To this end, we will discuss two assumptions that hold in specific scenarios. It is up to the practitioner to choose which assumption is more plausible for their setting (or come up with their own assumption along with updates to our formulae).

Assumption 1: Future is the past A simple assumption is to assume that observations from the past have not changed as of the time the study is being conducted. Formally, this means for each household, $X = X(t-1), Y = Y(t-1)$. Hence, (1) can be simplified as follows.

$$\begin{aligned} &P(Y_{X=x} = y^a \mid X(t-1) = x', Y(t-1) = y^p) \\ =&P(Y_{X=x} = y^a \mid X = x', Y = y^p) \qquad (3) \end{aligned}$$

This becomes the non-binary probability of necessity (Li and Pearl 2022a) of X on Y. Under this assumption, (3) can be bounded using the Theorem 7 in (Li and Pearl 2022a) (referred to as Li-Pearl's probability of necessity bounds).

Assumption 2: Constant pattern A relaxed assumption is to permit change in observations each year, but assume the changes follow the same pattern. Under this assumption, in addition to the observational data from the previous year $P(X(t-1), Y(t-1))$, we additionally need the observational data $P(X(t-2), Y(t-2))$ from the year before the previous year. Formally, this assumption translates to

$$\begin{aligned} &P(X(t-1) = x, Y(t-1) = y^p | X(t-2) = x', \\ &\qquad\qquad\qquad\qquad Y(t-2) = y^{pp}) \\ =\ &P(X = x, Y = y^p | X(t-1) = x', Y(t-1) = y^{pp}). \end{aligned}$$

Hence, we have the observational data $P(X, Y | X(t-1), Y(t-1))$ for (2). We can then use the observational data to bound the experimental probabilities to obtain $P(Y_X | X(t-1), Y(t-1))$. Given both observational and experimental data, we can use Li-Pearl's PNS bounds to bound (1).

Aspect 2: Unit Response Type Optimization of Subpopulations Without additional assumptions, counterfactual queries can rarely be point estimated, even with both observational and experimental data (Tian and Pearl 2000). However, it may be possible to sufficiently bound the query using observational and/or experimental data (Tian and Pearl 2000; Li and Pearl 2022a; Mueller, Li, and Pearl 2022; Zhang, Tian, and Bareinboim 2022; Dawid, Musio, and Murtas 2017; Li and Pearl 2022b). These existing works apply bounding methods to different settings, such as binary, continuous, monotonic, and other scenarios. We will adapt methods from (Li and Pearl 2022a) to bound (2) since the equation is a non-binary probability of causation. Specifically, we will apply Li and Pearl's Theorem 8, referred to as Li-Pearl's PNS bounds. In this setting, we will discuss the two assumptions above.

Assumption 1: Future is the past (2) can be simplified as follows.

$$\begin{aligned} &P(Y_{X=x} = y^a, Y_{X=x'} = y^b \mid X(t-1) = x', Y(t-1) = y^p) \\ =&P(Y_{X=x} = y^a, Y_{X=x'} = y^b \mid X = x', Y = y^p) \qquad (4) \end{aligned}$$

This is the non-binary probability of necessity and sufficiency (Li and Pearl 2022a) of X on Y. When $b = p$, (4) becomes the probability of necessity and can be bounded using Theorem 7 in (Li and Pearl 2022a). When $b \neq p$, (4) becomes 0.

Assumption 2: Constant pattern Under this assumption, similar to above, we have the observational data $P(X, Y \mid X(t-1), Y(t-1))$ for (2). We can then use the observational data to bound the experimental probabilities to obtain $P(Y_X \mid X(t-1), Y(t-1))$. Given both observational and experimental data, we can use Li-Pearl's PNS bounds to bound (2).

Computing the Benefit

Once we have the bounds of (2) (if assumption 2 holds) or (4) (if assumption 1 holds), there are multiple ways where the results can be used. For example, for each household, we can compute the bound of the expected EV mileage if added an additional EV and the bound of the expected EV mileage if not added an additional EV. So for each household, the difference in the two expectations is the expected EV mileage increment. We can identify which households are expected to have large mileage increments, and which are not. Another way the bounds of (2) or (4) can be used is to find what most likely to happen for each household, which means finding the PNS or PN bound with the highest probability. The practitioner can decide which way best fits their needs.

Aspect 1: Average Treatment Effect of Subpopulations The expected improvements for a subpopulation with $X(t-1) = x'$ and $Y(t-1) = y^p$ can be computed as follows.

$$\begin{aligned} &E[Y_{X=x} - Y_{X=x'} | X(t-1) = x', Y(t-1) = y^p] \\ &= \sum_y [P(Y_{X=x} = y | X(t-1) = x', Y(t-1) = y^p) - \\ &\qquad P(Y_{X=x'} = y | X(t-1) = x', Y(t-1) = y^p)] \cdot y \end{aligned}$$

Aspect 2: Unit Response Type Optimization of Subpopulations Given the counterfactual quantities, we can compute a generalized version (with higher dimensional outcome variables) of the benefit function (Li and Pearl 2019). In the setting of this paper, Y has 5 categories, hence there are $5 \cdot 5 = 25$ total *response types*. A household is of response type R_{ij} if the household's Y would be y^i if intervened with $X = x$ and Y would be y^j if intervened with $X = x'$, where i and j both have 5 possible values.

Let w_{ij} be the weight for the response type R_{ij}. The generalized expected benefit function for a subpopulation with $X(t-1) = x'$ and $Y(t-1) = y^p$ can be computed as follows.

$$\sum_{i,j \in \{1,\ldots,5\}} P[Y_{X=x} = y^i, Y_{X=x'} = y^j | X(t-1) = x', Y(t-1) = y^p] \cdot w_{ij}$$

We can also compute the *weighted probability of benefit*, which represents that for a subpopulation with $X(t-1) = x'$ and $Y(t-1) = y'$, what the probability is that an intervention of $X = x$ would result in larger Y compared to an intervention of $X = x'$.

$$P(\text{benefit}) = \sum_{1 \leqslant j < i \leqslant 5} P[Y_{X=x} > Y_{X=x'} | X(t-1) = x', Y(t-1) = y^p] \cdot w_{ij}$$

Experiment and Results

Aspect 2

Assumption 1: Future is the past In this section, we simulated the following example to illustrate our proposed framework under the first assumption.

We generated $P(X, Y, C)$ and $P(U)$ uniformly, as shown in Tables 1 and 2. We then applied Li-Pearl's causal effect bounds to derive the experimental (RCT) data $P(Y_x)$ using the data from Tables 1 and 2. The results are presented in Table 3. Subsequently, we used Li-Pearl's PN bounds to calculate the non-binary Probability of Necessity, with the results displayed in Table 4. Note that we are only presenting the upper bounds because the lower bounds for this randomly generated example are all zero, which provides no additional information.

Table 1: Scenario 1: Simulated observational distribution of the whole population.

	2 EV & 1 ICEV		**1 EV & 1 ICEV**	
	$\geq 15,000$	$< 15,000$	$\geq 15,000$	$< 15,000$
y^1	0.019	0.008	0.043	0.038
y^2	0.014	0.061	0.178	0.045
y^3	0.016	0.007	0.051	0.043
y^4	0.089	0.018	0.017	0.142
y^5	0.021	0.086	0.033	0.071

From Table 3, we have obtained narrow bounds on causal effects. However, according to our reasoning, the causal effects are not the correct queries we need. In Table 4, although there are only 10 entries that are not 1, we still gather useful information for decision making. For instance, we can determine that the probability of the population having 1 EV and 1 ICEV, with an EV miles level of y^2, and increasing their EV drive to level y^3 with one more EV, is at most 34.5%.

Table 2: Assumption 1: Prior knowledge about the trip type.

Typical trip type (U)	**Percentage**
More long trips (u)	6.7%
More short trips (u')	93.3%

Table 3: Assumption 1: Bounds of the experimental distribution.

	Lower bound	**Upper bound**
$P(y^1_x)$	0.027	0.117
$P(y^2_x)$	0.075	0.283
$P(y^3_x)$	0.023	0.100
$P(y^4_x)$	0.184	0.472
$P(y^5_x)$	0.164	0.438

Assumption 2: Constant pattern In this section, we simulated another example to illustrate our proposed framework under the second assumption. We generated $P(X, Y, C | X(t-1) = x', Y(t-1) = y^1)$ and $P(U | X(t-1) = x', Y(t-1) = y^1)$ uniformly, as shown in Tables 5 and 6. We then applied Li-Pearl's causal effect bounds to derive the experimental (RCT) data $P(Y_{X=x} | X(t-1) = x', Y(t-1) = y^1)$ using the data from Tables 5 and 6. The results are presented in Table 7. Subsequently, we used Li-Pearl's PNS bounds to calculate the non-binary Probability of Necessity and Sufficiency, with the results displayed in Table 8. Note that we are only presenting the upper bounds because the lower bounds for this randomly generated example are all zero, which provides no additional information.

From Table 7, we have obtained narrow bounds on causal effects. However, according to our reasoning, the causal effects are not the correct queries we need and should not be directly used to answer our question. In Table 8, for instance, we can determine that the probability of an individual household who has 1 EV and 1 ICEV, that would have EV miles level y^5 if they had 2 EV and 1 ICEV and would have EV miles level y^2 if remained with 1 EV and 1 ICEV, is at most 32.6%.

Conclusion

In this paper, we focus on the problem of optimizing electric vehicle procurement for maximizing environmental sustainability. We showed how to estimate how much each household benefits from an additional EV. This question is a causal question, and may need counterfactual reasoning, which is hard to solve using available observational or experimental data. To approach this problem, we developed a causal AI framework based on counterfactual tools. We showed how to apply this framework using a simulated experiment. For both assumptions discussed, we obtained bounds on the query of interest.

Table 4: Assumption 1: Upper bounds of non-binary Probability of Necessity.

$P(y_x^a \mid x', y^p)$	$p = 1$	$p = 2$	$p = 3$	$p = 4$	$p = 5$
$a = 1$	1	0.404	0.957	0.566	0.865
$a = 2$	1	0.933	1	1	1
$a = 3$	0.951	0.345	0.819	0.484	0.740
$a = 4$	1	1	1	1	1
$a = 5$	1	1	1	1	1

Table 5: Assumption 2: Simulated observational distribution of the population $X(t-1) = x', Y(t-1) = y^1$.

	2 EV & 1 ICEV		**1 EV & 1 ICEV**	
	$\geq 15,000$	$< 15,000$	$\geq 15,000$	$< 15,000$
y^1	0.110	0.062	0.066	0.035
y^2	0.066	0.009	0.030	0.072
y^3	0.142	0.047	0.006	0.063
y^4	0.010	0.098	0.029	0.098
y^5	0.004	0.013	0.012	0.028

Table 6: Assumption 2: Prior knowledge about the trip type.

Typical trip type (U)	**Percentage**
More long trips (u)	10.8%
More short trips (u')	89.2%

Table 7: Assumption 2: Bounds of the experimental distribution of the population $X(t-1) = x', Y(t-1) = y^1$.

	Lower bound	**Upper bound**
$P(y_x^1)$	0.173	0.393
$P(y_x^2)$	0.075	0.139
$P(y_x^3)$	0.211	0.395
$P(y_x^4)$	0.108	0.303
$P(y_x^5)$	0.017	0.051
$P(y_{x'}^1)$	0.101	0.476
$P(y_{x'}^2)$	0.102	0.394
$P(y_{x'}^3)$	0.069	0.260
$P(y_{x'}^4)$	0.127	0.521
$P(y_{x'}^5)$	0.040	0.188

Table 8: Assumption 2: Upper bounds of non-binary Probability of Necessity and Sufficiency of the population $X(t-1) = x', Y(t-1) = y^1$.

$P(y_x^a, y_{x'}^b)$	$b = 1$	$b = 2$	$b = 3$	$b = 4$	$b = 5$
$a = 1$	0.596	0.513	0.412	0.615	0.369
$a = 2$	0.439	0.356	0.255	0.458	0.212
$a = 3$	0.581	0.498	0.397	0.590	0.354
$a = 4$	0.570	0.487	0.386	0.589	0.343
$a = 5$	0.409	0.326	0.225	0.428	0.182

References

Burlig, F.; Bushnell, J.; Rapson, D.; and Wolfram, C. 2021. Low Energy: Estimating Electric Vehicle Electricity Use. *AEA Papers and Proceedings*, 111: 430–35.

Burnham, A.; Lu, Z.; Wang, M.; and Elgowainy, A. 2021. Regional emissions analysis of light-duty battery electric vehicles. *Atmosphere*, 12(11): 1482.

Dawid, P.; Musio, M.; and Murtas, R. 2017. The Probability of Causation. *Law, Probability and Risk*, (16): 163–179.

EPA. 2023. Inventory of U.S. Greenhouse Gas Emissions and Sinks: 1990-2021. Technical report, U.S. Environmental Protection Agency, EPA 430-R-23-002.

Jenn, A. 2020. Emissions benefits of electric vehicles in Uber and Lyft ride-hailing services. *Nature Energy*, 5(7): 520–525.

Li, A.; and Pearl, J. 2019. Unit Selection Based on Counterfactual Logic. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI-19*, 1793–1799. International Joint Conferences on Artificial Intelligence Organization.

Li, A.; and Pearl, J. 2022a. Probabilities of Causation with Non-binary Treatment and Effect. Technical Report R-516, Department of Computer Science, University of California, Los Angeles, CA.

Li, A.; and Pearl, J. 2022b. Unit selection with causal diagram. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, 5765–5772.

Mueller; and Pearl. 2022. Personalized Decision Making – A Conceptual Introduction. Technical Report R-513, Department of Computer Science, University of California, Los Angeles, CA.

Mueller, S.; Li, A.; and Pearl, J. 2022. Causes of effects: Learning individual responses from population data. In *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence (IJCAI-22)*, 2712–2718.

Nunes, A.; Woodley, L.; and Rossetti, P. 2022. Re-thinking procurement incentives for electric vehicles to achieve net-zero emissions. *Nature Sustainability*, 5(6): 527–532.

Pearl, J. 2009. *Causality*. Cambridge university press.

Tian, J.; and Pearl, J. 2000. Probabilities of causation: Bounds and identification. *Annals of Mathematics and Artificial Intelligence*, 28(1-4): 287–313.

Zhang, J.; Tian, J.; and Bareinboim, E. 2022. Partial counterfactual identification from observational and experimental data. In *International Conference on Machine Learning*, 26548–26558. PMLR.