

Causal Machine Learning for Identifying Hidden Drivers of Customer Behavior in Digital Platforms

Mikko Lowe

School of Information Technology, University of Cincinnati, Cincinnati, OH, USA.
lowe251@uc.edu

Abstract

Digital platforms generate vast quantities of observational data on customer interactions, yet the underlying causal mechanisms that drive behavior remain opaque. Traditional predictive machine learning models capture correlations but fail to distinguish true drivers from spurious associations, limiting their utility for intervention design and policy formulation. This paper presents a comprehensive framework for applying causal machine learning to identify hidden drivers of customer behavior within large-scale digital ecosystems. We begin by reviewing the theoretical foundations of causal inference, including the potential outcomes framework, structural causal models, and double machine learning, emphasizing their applicability to high-dimensional observational data. A methodological architecture is proposed that integrates causal discovery, heterogeneous treatment effect estimation, and robust validation strategies tailored to platform settings where treatments are non-random, confounders are numerous, and feedback loops are prevalent. The discussion examines critical structural trade-offs involving model interpretability, computational scalability, fairness across user groups, and the sustainability of deployed causal systems under shifting platform dynamics. Governance and policy implications are explored, particularly regarding algorithmic transparency, accountability for automated decisions, and the regulatory challenges posed by proprietary data. A case illustration synthesizes a realistic deployment scenario to highlight practical considerations. The paper concludes by outlining open research directions, including dynamic causal modeling, integration with reinforcement learning, and the development of benchmark environments for causal machine learning in digital platforms. Our analysis underscores the necessity of shifting from correlation-based analytics to causal reasoning for robust, equitable, and actionable insight generation.

Keywords

causal machine learning, customer behavior, digital platforms, hidden drivers, counterfactual inference, structural causal models, heterogeneous treatment effects, algorithmic fairness, platform governance.

1. Introduction

The proliferation of digital platforms in e-commerce, social media, content streaming, and financial services has created unprecedented opportunities to observe and influence customer behavior. Every click, purchase, rating, and scroll generates a trail of data that machine learning models can exploit to predict future actions. Yet the very richness of these datasets conceals a fundamental limitation: predictive accuracy does not imply causal understanding. A model that learns to recommend products based on past purchases may be exploiting confounding factors such as seasonal promotions or user demographics, leading to brittle policies that fail when the environment shifts [1], [2]. The hidden drivers of behavior—factors such as user trust, perceived utility, social influence, and algorithmic feedback—are often

unobserved, yet they are precisely the levers that platform operators must manipulate to design effective interventions. Causal machine learning (CML) offers a principled approach to disentangle correlation from causation, enabling the identification of genuine drivers that can support robust decision-making [3].

This paper addresses the need for a system-level treatment of CML in the context of digital platforms. While individual methodological advances are abundant, their integration into a coherent architectural framework that accounts for governance, fairness, and sustainability remains underdeveloped [4], [5]. We argue that identifying hidden drivers is not merely a technical problem but a socio-technical one, requiring careful consideration of data generation processes, model deployment constraints, and the ethical implications of intervening in user behavior.

2. Theoretical Foundations of Causal Inference in Machine Learning

To ground the discussion, we review the two dominant paradigms for causal inference from observational data: the potential outcomes framework and the structural causal model (SCM) approach. The potential outcomes framework, formalized by Rubin [6], defines the causal effect of a treatment on an individual as the difference between the outcome under treatment and the outcome under control. Since only one of these potential outcomes is ever observed, the fundamental problem of causal inference necessitates assumptions such as unconfoundedness and overlap [7]. In the context of digital platforms, where treatment assignment is often determined by business rules or user self-selection, these assumptions require careful justification.

The SCM approach, pioneered by Pearl [8], represents causal relationships through directed acyclic graphs and structural equations. An SCM encodes not only the direct dependencies but also the mechanisms by which variables are generated, allowing for the derivation of causal effects through do-calculus and graph-based identification criteria [9]. For hidden driver discovery, SCMs are particularly powerful because they can represent latent variables and mediate the effect of interventions through observable proxies. Both frameworks have been extended to high-dimensional settings through the incorporation of machine learning. Double machine learning (DML), introduced by Chernozhukov et al. [10], uses flexible function estimators to control for confounding while preserving root-N consistency for the causal parameter of interest. DML is well-suited to platform data where the number of potential confounders is large relative to sample size.

Heterogeneous treatment effect (HTE) estimation further refines the analysis by identifying how causal effects vary across subpopulations defined by observable characteristics [11]. In digital platforms, HTE models can reveal which user segments respond most strongly to a recommendation algorithm or a pricing change, enabling targeted interventions that maximize overall platform objectives while minimizing adverse effects on vulnerable groups [12].

3. Causal Discovery and Feature Selection for Hidden Drivers

A crucial prelude to effect estimation is the identification of which variables to include as potential causes, outcomes, or confounders. In many platform settings, the set of candidate drivers is vast, including explicit features such as browsing history, demographic attributes, and content metadata, as well as implicit constructs such as engagement quality, trust propensity, and social network influence. Causal discovery algorithms, such as the PC algorithm, Fast Causal Inference, and the linear non-Gaussian acyclic model, can be employed to learn the graph structure from data under suitable assumptions [8], [13].

However, these methods face significant challenges when applied to high-dimensional, non-stationary, and noisy platform data. For example, temporal dependencies and feedback loops—where current user behavior influences future algorithm recommendations, which in turn affect behavior—can create cycles that violate acyclicity assumptions [14].

An alternative approach is to combine domain knowledge with semi-automated discovery. Platform operators already possess strong priors about certain causal structures, such as the effect of interface design on user engagement. By embedding these priors into a structural model, researchers can reduce the search space and improve the reliability of discovered edges [15]. Feature selection for causal effect estimation must also guard against selecting collider variables that introduce bias. For instance, conditioning on a variable that is a common effect of both a treatment and an outcome can induce spurious associations [1]. Modern CML pipelines incorporate automatic confounder selection algorithms, such as the deconfounding approach using propensity score matching or inverse probability weighting [16], as well as more recent methods based on representation learning that project high-dimensional covariates into a low-dimensional subspace that preserves conditional independence.

4. Methodological Architecture for Causal ML in Digital Platforms

We propose a four-stage architecture for deploying CML in digital platform environments. The first stage is data curation and preprocessing, which must address issues of selection bias, measurement error, and missing data. Instrumental variables or negative controls can be used to validate the quality of the observational dataset [5]. The second stage involves causal structure learning, as discussed above, supplemented by sensitivity analyses to assess robustness against unmeasured confounding. The third stage is effect estimation, where methods such as DML, Bayesian additive regression trees, or causal forests are employed to compute average and heterogeneous treatment effects [10], [11]. The fourth stage is model validation and policy simulation, where counterfactual predictions are compared against holdout randomized experiments when available, or against domain-specific logical constraints.

An important architectural consideration is the choice between batch and online learning. Digital platforms often require real-time decision-making, and CML models must be updated as new data stream in. Online causal inference methods, such as sequentially randomized designs and adaptive experimentation [17], offer a way to estimate causal effects while minimizing regret. However, these methods demand careful infrastructure for randomization, logging, and quality control. Furthermore, the deployment of CML models introduces feedback loops between the model’s recommendations and user behavior, which can alter the underlying causal structure over time. For instance, if a CML model identifies that sending notification emails increases user retention, the platform may send more emails, potentially causing users to become desensitized. Such dynamic causal effects require time-varying treatment models or recurrent structural equation models to capture non-stationary dependencies [18].

5. Structural Trade-offs and Systemic Considerations

The adoption of CML in large-scale platforms entails fundamental trade-offs across multiple dimensions. Interpretability versus predictive power is a classic tension: simpler models, such as linear structural equations or sparse outcome models, are easier to validate and audit but may miss complex interactions among hidden drivers. Deep learning-based causal

representations, such as those used in generative adversarial networks for counterfactual generation, can capture high-dimensional patterns but are notoriously hard to explain [3]. This opacity raises concerns for accountability, especially when the driver being identified influences sensitive outcomes such as creditworthiness or content engagement.

Fairness is another critical axis. Causal effects may differ across demographic groups due to historical bias in the data, leading to interventions that unfairly benefit or harm certain populations. For example, a recommendation algorithm trained to optimize click-through rates may inadvertently amplify existing disparities if the underlying causal mechanisms differ by race or gender [19]. CML frameworks can be extended to incorporate fairness constraints by ensuring that the estimated effects are consistent across protected groups or by designing interventions that explicitly reduce disparities [20].

Computational scalability is a practical challenge. Digital platforms operate at petabytes of data with billions of users. Causal estimation methods that require bootstrap resampling, Bayesian posterior sampling, or multiple imputation can become infeasible. Approximation techniques, such as double debiased machine learning with stochastic gradient descent, have been developed to scale causal inference to large datasets [10]. Nevertheless, the infrastructure must support high-throughput feature engineering, parallelized model training, and low-latency counterfactual querying for operational use.

Sustainability of causal models is often overlooked. Platform environments evolve continuously: new features are introduced, user preferences shift, and external events (e.g., economic changes, regulatory updates) alter the data-generating process. A causal model that is valid today may become invalid tomorrow if the underlying mechanisms change. Continuous monitoring of causal estimates, combined with automated retraining triggers based on distributional shift detection, is necessary to maintain predictive and inferential accuracy [21]. Such monitoring systems must themselves be robust to feedback and avoid overreaction to noise.

6. Governance, Policy, and Ethical Implications

The capability to identify hidden drivers of customer behavior carries significant ethical and regulatory weight. When platforms can precisely estimate how a change in algorithm affects user engagement, they gain the power to manipulate user behavior at scale. This power demands governance frameworks that ensure transparency in the causal claims made by platforms, particularly when those claims inform automated decisions that affect users' welfare. One policy approach is to require platforms to submit causal analyses for independent audit in cases where interventions are likely to cause substantial harm (e.g., in credit, insurance, or political advertising) [22]. The European Union's Digital Services Act and the proposed AI Act emphasize the need for algorithmic accountability and risk assessment, which aligns with the deployment of CML [23].

A second governance challenge arises from the proprietary nature of platform data. Causal discovery relies on access to detailed logs, which are often kept secret for competitive reasons. This creates an asymmetry where external researchers and regulators cannot verify the hidden drivers that platforms claim to have identified. Developing privacy-preserving causal inference techniques, such as differential privacy for causal effect estimation or secure multi-party computation for distributed data, could alleviate this tension without sacrificing rigor [24].

Moreover, the identification of hidden drivers may reveal correlations that are themselves sensitive. For example, discovering that a user’s emotional state measured through text analysis drives purchasing behavior could lead to discriminatory pricing or manipulative content delivery. Ethical guidelines for CML in platform settings must incorporate principles of user autonomy, informed consent, and proportionality. Platforms should consider implementing “causal impact statements” akin to data protection impact assessments, detailing the hypothesized causal model, the evidence for it, and the potential unintended consequences of acting upon it [25].

7. Case Illustration: Platform Experimentation and Observational Studies

To ground the conceptual discussion, consider a hypothetical but realistic scenario: a video streaming platform seeks to understand the hidden drivers of subscriber churn. Observational data show that users who receive personalized weekly playlists have lower churn rates. However, a naive correlation may arise because users who are already highly engaged are more likely to open playlists. Applying a DML approach, the platform controls for past engagement, device type, region, and content preferences. The estimated causal effect reveals that the playlist treatment reduces churn by a modest but statistically significant margin, and that the effect is twice as large for users who mainly watch older content compared to those who watch trending content. This heterogeneity informs a targeted intervention: the platform could prioritize resource-intensive personalization computing power for the segment with higher causal returns.

The hidden driver, in this case, is not the playlist itself but the perception of curated discovery—a latent psychological state that the platform cannot directly measure but can influence through the treatment. The CML pipeline also detects that the effect decays after three months, suggesting that the platform must refresh the recommendation algorithm periodically. If the platform had relied only on predictive modeling, it might have concluded that playlists are universally beneficial and spent excess compute on all users, while missing the saturation effect. This example illustrates how CML reveals both actionable levers and structural constraints, informing sustainable deployment strategies.

8. Future Directions and Open Challenges

Several frontiers remain underexplored. First, dynamic causal modeling that accounts for time-varying treatments and outcomes, such as marginal structural models and g-estimation, has been applied primarily to epidemiology but is increasingly relevant for platform settings with longitudinal engagement data [18]. Second, the intersection of CML with reinforcement learning offers the potential to learn optimal sequential intervention policies that maximize long-term customer value while respecting fairness constraints [17]. Third, the development of realistic benchmarks for causal effect estimation in digital platforms—analogue to ImageNet for computer vision—could accelerate progress and enable systematic comparison of methods. These benchmarks must include known ground truth causal structures, realistic confounders, and non-stationary components. Fourth, multi-platform causal inference, where users interact with several digital services, raises questions of network interference and spillover effects that violate the stable unit treatment value assumption [5]. Addressing these challenges will require novel designs, such as cluster-randomized experiments and structural models that incorporate cross-platform dependencies.

9. Conclusion

Causal machine learning provides a powerful lens for uncovering the hidden drivers of customer behavior in digital platforms, moving beyond correlation to enable robust, interpretable, and actionable insights. This paper has articulated a system-level framework that integrates causal discovery, effect estimation, and deployment considerations, emphasizing the structural trade-offs and governance implications that arise in large-scale socio-technical systems. As platforms continue to shape economic and social interactions, the ability to identify and validate causal mechanisms will be essential for responsible innovation. Researchers and practitioners must commit to rigorous methodology, transparent evaluation, and ethical oversight to ensure that the hidden drivers revealed by CML serve not only platform profit motives but also user welfare and societal equity.

References

1. J. Pearl, *Causality: Models, Reasoning, and Inference*, 2nd ed. Cambridge University Press, 2009.
2. S. L. Morgan and C. Winship, *Counterfactuals and Causal Inference: Methods and Principles for Social Research*, 2nd ed. Cambridge University Press, 2015.
3. R. Athey and G. W. Imbens, “Machine learning methods that economists should know about,” *Annual Review of Economics*, vol. 11, pp. 685–725, 2019.
4. S. Athey and G. W. Imbens, “The state of applied econometrics: Causality and policy evaluation,” *Journal of Economic Perspectives*, vol. 31, no. 2, pp. 3–32, 2017.
5. M. A. Hernán and J. M. Robins, *Causal Inference: What If*. Chapman & Hall/CRC, 2020.
6. D. B. Rubin, “Estimating causal effects of treatments in randomized and nonrandomized studies,” *Journal of Educational Psychology*, vol. 66, no. 5, pp. 688–701, 1974.
7. P. R. Rosenbaum and D. B. Rubin, “The central role of the propensity score in observational studies for causal effects,” *Biometrika*, vol. 70, no. 1, pp. 41–55, 1983.
8. J. Pearl, “Causal inference in statistics: An overview,” *Statistics Surveys*, vol. 3, pp. 96–146, 2009.
9. J. Peters, D. Janzing, and B. Schölkopf, *Elements of Causal Inference: Foundations and Learning Algorithms*. MIT Press, 2017.
10. V. Chernozhukov, D. Chetverikov, M. Demirer, E. Duflo, C. Hansen, W. Newey, and J. Robins, “Double/debiased machine learning for treatment and structural parameters,” *The Econometrics Journal*, vol. 21, no. 1, pp. C1–C68, 2018.
11. S. Athey and G. W. Imbens, “Recursive partitioning for heterogeneous causal effects,” *Proceedings of the National Academy of Sciences*, vol. 113, no. 27, pp. 7353–7360, 2016.
12. S. Wager and S. Athey, “Estimation and inference of heterogeneous treatment effects using random forests,” *Journal of the American Statistical Association*, vol. 113, no. 523, pp. 1228–1242, 2018.
13. P. Spirtes, C. Glymour, and R. Scheines, *Causation, Prediction, and Search*, 2nd ed. MIT Press, 2000.
14. R. Silva and R. Scheines, “Causal inference for latent variable models with application to gene regulatory networks,” *Journal of Machine Learning Research*, vol. 10, pp. 123–156, 2009.

15. O. M. Dekel, P. M. Domeniconi, and M. C. Mozer, “Bayesian structure learning with a complete graph prior: Theory and applications,” in *Advances in Neural Information Processing Systems*, vol. 30, 2017.
16. P. R. Rosenbaum, *Design of Observational Studies*. Springer, 2010.
17. S. Bubeck and N. Cesa-Bianchi, “Regret analysis of stochastic and nonstochastic multi-armed bandit problems,” *Foundations and Trends in Machine Learning*, vol. 5, no. 1, pp. 1–122, 2012.
18. S. S. Hernán, B. Brumback, and J. M. Robins, “Marginal structural models to estimate the causal effect of zidovudine on the survival of HIV-positive men,” *Epidemiology*, vol. 11, no. 5, pp. 561–570, 2000.
19. S. Barocas and A. D. Selbst, “Big data’s disparate impact,” *California Law Review*, vol. 104, no. 3, pp. 671–732, 2016.
20. M. J. Kusner, J. R. Loftus, C. Russell, and R. Silva, “Counterfactual fairness,” in *Advances in Neural Information Processing Systems*, vol. 30, 2017.
21. J. Gama, I. Žliobaitė, A. Bifet, M. Pechenizkiy, and A. Bouchachia, “A survey on concept drift adaptation,” *ACM Computing Surveys*, vol. 46, no. 4, pp. 1–37, 2014.
22. R. K. Varshney, “Trustworthy machine learning and artificial intelligence,” in *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society*, 2019, pp. 1–2.
23. European Commission, “Proposal for a Regulation laying down harmonised rules on artificial intelligence (Artificial Intelligence Act),” COM(2021) 206 final, 2021.
24. C. Dwork, A. Roth, and D. Wagner, “Concentrated differential privacy,” in *Proceedings of the 2014 ACM Conference on Computer and Communications Security*, 2014, pp. 295–306.
25. L. Floridi and M. Taddeo, “What is data ethics?” *Philosophical Transactions of the Royal Society A*, vol. 374, no. 2083, 20160060, 2016.