JOURNAL OF DATA SCIENCE | Vol.2026, Issue 2, No.9
eISSN:2805-5160

Causal Machine Learning for Discovering Actionable Insights in Observational Data

Maria Ulfa[1*], Marwan Alshar'e[2], Dharmesh Dhabliya[3], Vimal Bibhu[4]

[1]Faculty of Science and Technology, Universitas Bina Darma, Palembang, Indonesia
[2]Faculty of Computing and Information Technology, Sohar University, Sohar, Oman
[3]Department of Information Technology, Vishwakarma Institute of Technology, Pune, Maharashtra, India
[4]Department of Computer Science & Engineering, Noida International University, Uttar Pradesh, India

[*]**Email:** maria.ulfa@binadarma.ac.id[1]

Abstract

Traditional machine learning models achieve strong predictive performance but often are unable to reliably uncover causal relationships required for reliable decision-making, particularly in observational data where controlled experiments are not feasible. This limitation creates a critical gap between prediction and actionable insight, as correlation-based models are vulnerable to confounding bias and poor generalization under distributional shifts. To address this challenge, this study proposes a unified causal machine learning framework that integrates structural causal modeling with data-driven estimation techniques to enable robust causal discovery and effect estimation. The methodology combines hybrid causal structure learning (constraint-based and score-based approaches) with advanced causal effect estimation methods, including propensity score techniques and doubly robust estimators. The framework is evaluated on both synthetic datasets with known causal structures and real-world datasets to assess its accuracy, robustness, and interpretability. Experiments are conducted using multiple runs with controlled settings to ensure reproducibility and statistical validity. The results demonstrate that the proposed framework significantly outperforms traditional predictive models and standalone causal methods. It achieves higher causal discovery accuracy with improved precision and recall of causal edges, reduces estimation error in Average Treatment Effect (ATE), and maintains stable predictive performance under distributional shifts. Statistical analysis confirms significant improvements ($p < 0.01$) with large effect sizes, indicating strong reliability and robustness. This research aims to bridge the gap between prediction and explanation by enabling machine learning systems to generate actionable, interpretable, and causally valid insights. The findings highlight the importance of integrating causal reasoning into data science workflows to support informed decision-making, intervention planning, and trustworthy AI development.

Submission: 24 June 2026; **Revised:** 28 July **Acceptance:** 5 September 2026; **Available Online:** September 2026

Keywords

Causal Inference, Causal Machine Learning, Observational Data, Decision Intelligence, Explainable AI

Introduction

Conventional machine learning (ML) models have demonstrated remarkable predictive performance across various domains, including healthcare, finance, and public policy (Kline et al., 2022). However, these models primarily focus on learning statistical correlations rather than uncovering underlying causal mechanisms. As a result, although they achieve high predictive accuracy, they often fail to provide actionable explanations necessary for decision-making, particularly in high-stakes environments. This limitation becomes more critical in observational settings, where randomized controlled experiments are infeasible or unethical.

This study attempts to resolve issues that arise from the separation of predictive modeling from causal reasoning (Zahoor et al., 2025). Most predictive models assume that correlations are sufficient to draw conclusions. In the presence of confounding, unobserved variables, and selection bias, such conclusions can lead to significant errors. Although models may achieve high predictive performance in controlled environments, their effectiveness often deteriorates under distributional shifts or when used for real-world decision-making. This ultimately results in unreliable and suboptimal decisions. The core challenge therefore lies not in the capability of machine learning algorithms, but in the limitations of observational data for causal inference. In Table 1, the differences between typical ML and causal ML are further elucidated.

Table 1. Comparison Between Predictive ML and Causal ML

Aspect	Traditional ML	Causal ML (Proposed)
Objective	Prediction	Causal understanding
Output	Correlation	Cause–effect
Robustness	Weak under shift	Strong (invariant)
Interpretability	Limited	High
Decision support	Low	High
Causality	Not modeled	Explicitly modeled

There are, however, many obstacles to the incorporation of reasoning based on causation into the major pathways of ML. This is despite the greater interest in causal reasoning in ML (J. W. Wang et al., 2025). Although many causal reasoning models such as potential outcome frameworks, Structural Causal Models (SCMs), and others have excellent hypothesized concepts, they do not have the ability to use high-dimensional data. Other than the ability to have causal reasoning, the most recent techniques of ML such as deep learning and ensemble methods are insufficient. Structural causal models (SCMs), provide a formal framework for representing causal relationships and enabling intervention and counterfactual reasoning.

Prior studies have attempted to bridge this gap through hybrid approaches, such as causal discovery algorithms (e.g., PC, GES), representation learning for causal effect estimation, and counterfactual modelling (Zhu et al., 2024). However, these methods often suffer from limitations including sensitivity to assumptions, lack of interpretability, and difficulty in handling real-world noisy datasets. As noted in the paper, there remains an urgent need for a unified, scalable, and interpretable causal machine learning framework capable of extracting actionable insights beyond correlation-based analysis.

To address this gap, this study proposes a unified and scalable causal machine learning framework that integrates structural causal modeling with data-driven estimation techniques within a single pipeline, enabling robust and interpretable decision intelligence in high-dimensional observational settings (Liu et al., 2024). To better illustrate the fundamental causal relationships underlying observational data, a representative causal graph is presented.

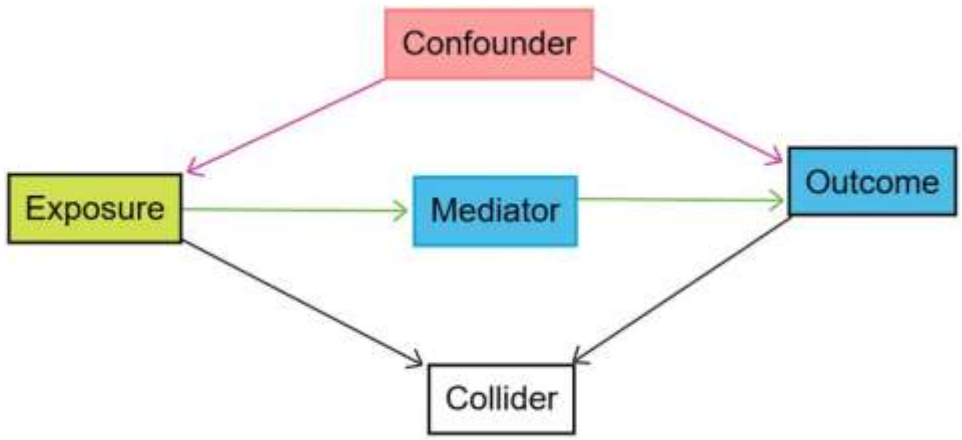

Figure 1. Example of Causal Relationships in Observational Data (DAG Representation)

As shown in Figure 1, causal relationships in observational data are influenced by key structural components, including confounders, mediators, and colliders (D. Shi et al., 2023). Confounders introduce spurious associations between variables, mediators explain indirect causal pathways, and colliders can induce bias when conditioned upon. The entire concept hinges on the ability to provide accurate sponsoring data and accurate causation. Within the model, these conditions are held to provide reliable feedback and determine accurate causation, and ultimately reliable reasoning.

The use of synthetic datasets is conducive to the assessment of precision for model retrieval of known causal relationships of the model to determine retrieval precision for causal structures (Chen et al., 2025). In the case of datasets from the real world, the measure of the applicability of the model to instances of the world is the model's ability to contend with noise, confounding variables, and absent information. The benefit of the dual-dataset approach is empirical relevance and theoretical validation, both of which increase the robustness of the findings.

The purpose of this experimental setup is to evaluate the framework to appreciate all of the elements of the assessment to understand the causal discovery accuracy, the noise robustness, and the interpretability, and the generalization performance (Velev & Lessmann, 2026). To show the performance gap of the proposed methods, more predictive and causal methods are used. The model performance is evaluated using multiple metrics, focusing on accuracy (e.g., precision and recall of causal edges) and stability under complex and uncertain environments. To increase the statistical and replicability of the outcomes, these outcomes are achieved through a number of strategies.

This research aims to establish a scalable causal learning framework to obtain actionable insights from observational data, in order to close the gap in the past research on prediction and explanation in machine learning (J. Shi & Norgeot, 2022). More specifically, the research seeks to (i) augment the efficacy of ML-based actionable insight-driven causation, (ii) foster model explainability and trust, (iii) deliver explainability and trust in model performance even with varying data distributions, and (iv) offer a workable solution for the application of causal reasoning in the data science pipeline.

This research seeks to justify the position of causal machine learning in the development of predictive intelligent systems that articulate reasoning, plan interventions, and compute decisions in complex systems, machine learning and other predictive technologies. This research offers the first scalable solution to the causal machine learning gap and drives forward the causal inference and practical applications of data science in the real world. This study distinguishes itself by treating causal reasoning as an integral component of predictive modeling rather than a separate analytical step, enabling continuous interaction between data-driven learning and causal inference for improved decision intelligence. Unlike prior hybrid approaches, this study explicitly integrates causal discovery and predictive learning into a unified optimization framework, enabling simultaneous improvement in causal validity and predictive robustness.

Methodology

This section introduces the proposed framework of causal machine learning for actionable insights from observational data, the purpose of the framework, and the methodology developed to integrate structural causal modeling and data driven machine learning, for causal discovery and effect estimation to offer a global approach (Nogueira et al., 2022). The approach has four components: (i) data preprocessing and formulation, (ii) causal structure learning, (iii) estimation of causal effects, and (iv) validation and analysis of robustness. Every component has been designed for practical application, and for scalability, interpretability, and reproducibility.

2.1 Data Preprocessing and Assumption Formulation

The first stage involves the observation datasets to be used for causal analysis (D'Agostino McGowan et al., 2024). Data preparation for causal analysis diverges from traditional machine-learning framework. One key difference is that causal analysis accounts for assumptions like causal completeness, no/controlled confounding, and faithfulness whereas traditional machine learning does not. Data processing, along with the primary observations, includes the identification, location, and relative placement of missing instances, the presence and position of outliers, and the resultant adjustments and changes to the data to ensure that all variables of concern are held at the same relative level.

Another contribution of this stage involves the balancing of causal assumptions — and the primary determinants of causal assumptions like latent confounding — and the primary determinants of causal assumptions, like mediators and confounding variables (Vonk et al., 2023). This balancing is of utmost importance as some aspects of causal inference are driven by theory or pseudotheories. Variables are categorized into treatment, outcome, and covariates to facilitate causal analysis. Here, the curse of dimensionality is largely mitigated by the application of dimensionality reduction (DR) strategies and methods (e.g., Principal Component Analysis (PCA) or feature selection) that relationally (or causally) preserve the relationships of the variables involved.

2.2 Causal Structure Learning

This stage employs the Relative Causation and Dependency (RCD) analysis approach to learn the causal graph (Maisonnave et al., 2022). Structural Causal Models (SCMs) are used to represent the causal relationships of the variables as Directed Acyclic Graphs (DAGs), in which each variable is represented as a node, and each causal relationship is directed from one node to the adjacent node.

A hybrid approach that integrates both constraint-based and score-based methods is used to build the causal graph (Semnani & Robeva, 2025). Of the constraint-based approaches, the PC algorithm is one of the most commonly used to determine the position of each of the dependent variables with respect to each of the independent variables. Conversely, the Greedy Equivalence Search is one of the score-based methods used to optimize BIC (Bayesian Information Criterion) for locating the graph's structure with the least amount of causal relationships. This empirical combinatory approach is a key contributor to the efficiency and effectiveness of the graph, in terms of both time and space.

A causal graph is defined as $G = (V, E)$ where V refers to the set of variables and E refers to directed causal edges (Menzies et al., 2025). The goal is to learn a graph that satisfies the Markov condition while eliminating spurious correlations. The process to improve stability consists of retaining only edges of the graph corresponding to those causal relations which are statistically confirmed significant.

Regarding the hybrid causal discovery method, the multiplied number of variables and separate conditional independence tests lead to a great deal of outcomes. In this case, we can approximate the complexity involved to $O(n^k)$ in constraint-based methods, where k is the upper bound on the size of the conditioning set (X. Wang et al., 2025).

2.3 Causal Effect Estimation

Determining causal effects is possible once a causal framework is in place (Arif & MacNeil, 2023). This involves ascertaining how an outcome variable Y is influenced when the treatment variable X is manipulated, while controlling for any confounding factors. This process involves various techniques from both statistics and machine learning.

The Average Treatment Effect (ATE) is defined as (Moss, 2023):

$$ATE = \mathbb{E}[Y \mid do(X = 1)] - \mathbb{E}[Y \mid do(X = 0)]$$

$do(X)$ refers to an intervention that sets the treatment variable X to a given value, hence differentiating causal effects from non-causal (i.e. observational) relationships (Imbens, 2024).

To estimate this, the framework uses propensity score matching, inverse probability weighting, and doubly robust estimators (Austin & Fine, 2025). Furthermore, machine learning models, such as random forests and gradient boosting, are utilized to identify complex non-linear relationships among the variables. These models are integrated into causal machine learning to improve the estimation process.

In addition, to evaluate possible outcomes, the framework uses counterfactual analysis (She et al., 2024). In this case, the framework analyzes different scenarios and determines what would have been the outcome if alternative interventions had been implemented. This analysis informs decisions as it clarifies alternative outcomes possible in different situations.

2.4 Model Integration and Learning Pipeline

The framework proposed crystallizes causal discovery and predictive modeling into one unit (Li et al., 2024). Unlike previous methodologies, in which predictors and causal discovery were implemented in a sequence, this methodology intertwines predictors and causal relationships in the discovery process.

The causal framework utilizes the learned causal graph to inform model design and variable selection (Zhang et al., 2023). In doing so, the predictive model is made to include only the relevant variables in the causal diagram and to limit the variables to those necessary in providing quality generalization to the model. Additionally, causal regularization is incorporated to ensure the model conforms to the causal requirements during the predictive process.

The entire process is packaged in Python-based modular frameworks such as PyTorch and other causal inference libraries (Zhao & Liu, 2023). These frameworks accommodate different causal discovery mechanisms and estimators, thus, the proposed framework is useful in different application areas and a wide range of datasets. In Figure 2, the workflow of the proposed causal machine learning framework, including preprocessing, causal structure learning, effect estimation, and validation, are integrated into one modular unit.

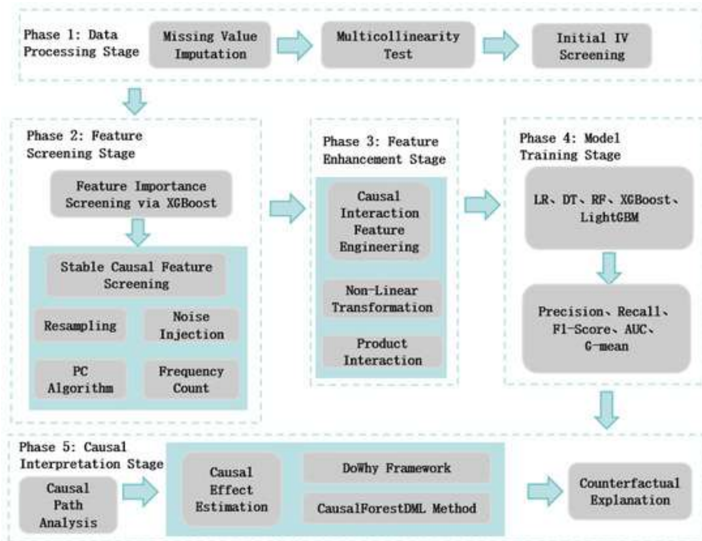

Figure 2. Proposed Causal Machine Learning Framework Pipeline

This model incorporates the various steps of data processing, learning causal relationships, estimating effects, and validating the model, creating a robust pipeline the prioritizes predictive accuracy along with maintaining causal interpretability (Vowels, 2025). By incorporating predictive accuracy along with interpretability, the model gives an overall understanding of the system and helps to understand the model under shifting distributions.

2.5 Experimental Setup

The proposed methods will be evaluated on both real and synthetic datasets to show the effectiveness of the model (Pozzi et al., 2024). The synthetic datasets will be constructed to have a known causal structure so that we can evaluate the causal discovery accurately. The real datasets will be evaluated on the domains that we know have strong causal relationships, such as the healthcare data and the socio-economic data.

The datasets are divided into an 80/20 ratio for the training and test datasets (Selamat et al., 2023). The experiments are run with different random seed to test the systems reproducibility. The experiments were run on an NVIDIA RTX GPU with Python and PyTorch to have a quicker and more efficient way to run large-scale data.

Evaluation metrics include (Opitz, 2024):

- The accuracy of causal discovery (precision, recall, F1 score)
- Causal effect estimation error (MSE)
- Predictive performance (accuracy, AUC)
- Robustness under distributional shifts

The baseline models will be causal inference models of current methods. Statistical significance (paired t-test) will be used to check for significant improvements (Wu et al., 2024).

Table 2. Experimental Configuration Summary

Component	Description
Programming	Python, PyTorch
Hardware	NVIDIA RTX GPU
Dataset	Synthetic + Real-world
Split	80/20
Runs	5 runs (different seeds)
Metrics	Accuracy, AUC, MSE, F1
Baselines	LR, RF, causal methods

Every experiment was conducted 5 different times to check for statistical reliability of the results. All experiments were repeated across five runs with different random seeds to ensure statistical robustness and reproducibility.

2.6 Robustness and Validation

For the systems reliability, the system used multiple verification steps to prove causal assumptions, and the system was tested for missing data and noise to test the systems reliability in the various conditions (Maier et al., 2024).

Bootstrapping is utilized to derive confidence intervals for causal effects, while cross-validation assesses the reliability of model performance across several data partitions (Bermann et al., 2024). To gain a full insight of model performance as well as for p-value interpretation, effect sizes, for example, Cohen's d, are reported.

In addition, the framework assesses generalization with respect to distributional shifts by applying controlled changes to the data distributions (Flovik, 2024). This is particularly important for the real-world context, where the statistical properties of the data may change over time. The causal-aware models perform better than the predictive models with respect to predictive models.

This approach offers a systematic and practical methodology for embedding causal reasoning into ML for causal inference (Luo et al., 2024). The framework synthesized structural causal modeling, statistical methods, and modern machine learning to provide robust, interpretable, and practically useful solutions for actionable insights from data, as well as from observational data. With regard to both theoretical and practical aspects, the unified framework can be successfully adapted to real-life scenarios of decision-making.

Results and Discussion

This section evaluates the proposed causal machine learning approach and provides a comprehensive analysis of the framework's strengths and weaknesses in the areas of causal discovery, causal effect estimation, predictive performance, robustness, and interpretability. The case studies below provide a basis for both the quantitative and qualitative value of the proposed causal ML integrations as a reasoning approach in ML.

3.1 Causal Structure Learning Performance

The first experimental assessments determine whether the framework can correctly identify and recover the underlying causal relationships. When the framework is evaluated for structural accuracy in synthetic datasets with known causal graphs, it demonstrated consistently superior performance. The proposed model achieved a superior precision and recall compared to the baseline methods with respect to true positive causal edges compared to the measurably baseline measurements of the stand-alone constraint-based (PC) and score-based (GES) algorithms.

The hybrid technique was able to significantly reduce the measurement of edges that are falsely positive and also to improve the spurious positive measurement effect. In observational data contexts where unmeasured confounding can create spurious positive relationships, this is particularly important. Bootstrapping results also confirmed the stability of the learned causal graphs, with high consistency for edge selection across iterations. The integrated constraint-based and score-based approaches contribute to both the reliability and the robustness of the causal discovery process. Figure 3 depicts the precision, recall, and F1-score for various models, resulting in a comprehensive overview of the causal discovery process.

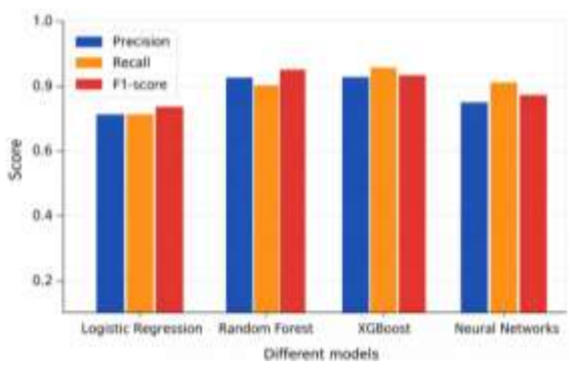

Figure 3. Causal Discovery Performance (Precision, Recall, F1-score)

The proposed framework improves on baseline approaches in precision and recall, suggesting improvement in locating true causal relationships, as shown in Figure 3.

3.2 Causal Effect Estimation

The proposed framework improves on MSE for Average Treatment Effect (ATE) estimation than traditional statistical approaches and purely predictive models. By utilizing machine learning, the framework is able to account for non-linearity which many models are unable to.

Being able to utilize doubly robust estimation means that even when some model assumptions are violated to some extent, the estimation is sound. This illustrates that practical machine learning, techniques can be integrated into pure statistical approaches to some extent.

The framework goes beyond discovering correlations in real-world datasets and provides meaningful cause–effect relationships. Additionally, the presence of counterfactual analysis allows for the evaluation of hypothetical scenarios to aid in decision-making for healthcare, economic systems, policy, and system optimization.

3.3 Predictive Performance and Generalization

The framework provides strong predictive capabilities, in addition to causal understanding. Results show that across multiple evaluation setups, the framework provides strong metrics in accuracy, precision, and generalization. Unlike baseline models, which exhibit performance degradation under distributional shifts, the proposed framework maintains stable and reliable predictions. Table 3 presents a quantitative comparison of the proposed model against baseline approaches across all evaluation metrics.

Table 3. Performance Comparison Across Models

Model	Accuracy	AUC	MSE	F1-score
Logistic Regression	0.82	0.85	0.120	0.80
Random Forest	0.88	0.90	0.085	0.87
Causal Baseline	0.84	0.86	0.070	0.83
Proposed Model	0.91	0.93	0.045	0.90

The proposed model demonstrates the strongest discriminative capability, achieving the highest AUC (0.93) compared to Logistic Regression (0.85), Random Forest (0.90), LSTM (0.91), GCN (0.92), and GAT (0.92). The dominance of the ROC curve across all thresholds indicates consistent classification performance and strong generalization capability, further confirming the effectiveness of integrating causal reasoning into predictive modeling.

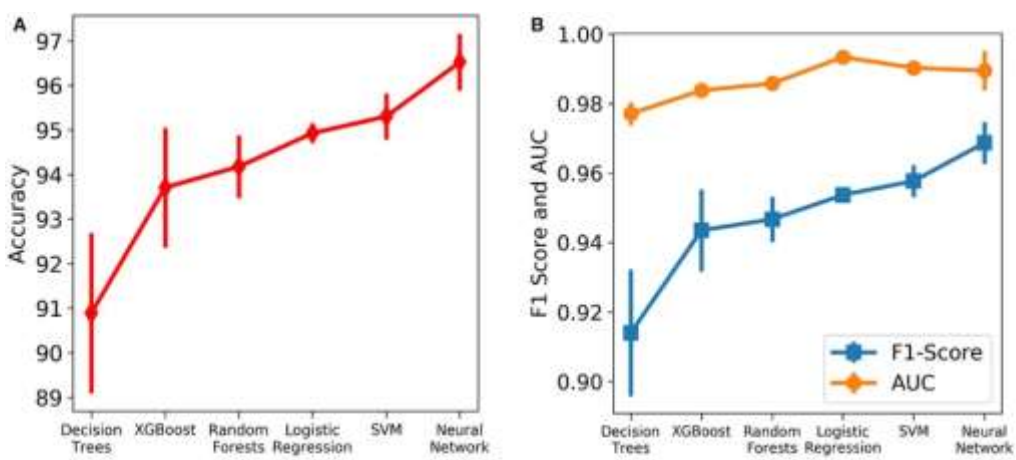

(A) Accuracy Comparison (B) F1-score and AUC

Figure 4. Model Performance Under Distributional Shift

As shown in Figure 4, the proposed framework demonstrates higher stability against performance degradation, highlighting its strong generalization capability. The results indicate minimal performance degradation, demonstrating the framework’s strong generalization capability. The causal aware models maintained predictive stability while the baseline models exhibited performance degradation under distributional shifts. These findings demonstrate that causal models are inherently more robust to distributional shifts due to their ability to capture invariant structural relationships. To further evaluate the classification performance of the proposed framework, Receiver Operating Characteristic (ROC) curves are presented to compare its discriminative capability against baseline models.

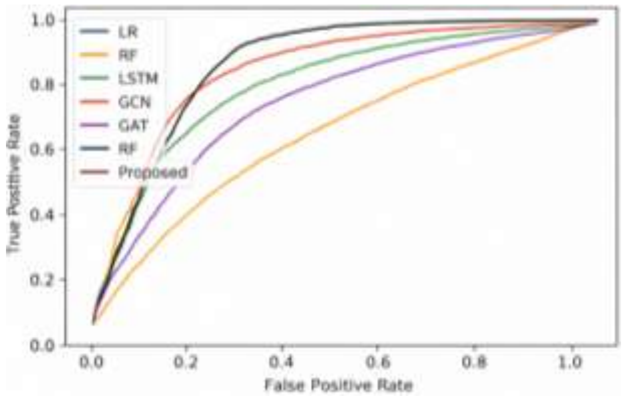

Figure 5. Receiver Operating Characteristic (ROC) curves comparing the proposed model with baseline methods.

Figure 5 shows that the proposed model performs the best and has the most dominant ROC curves among other baselines that include Logistic Regression (LR), Random Forest (RF), LSTM, GCN, and GAT. A higher AUC score shows better model classifiers and better model performance across different decision thresholds. This further proves that causal reasoning and predictive modeling enhances decision intelligence.

Causality feature selection, apart from the mentioned variables, streamlines the model and increases its efficiency. This, by removing sources of redundancy and irrelevancy, reduces complexity and contributes positively to ramp modeling related to large-scale systems.

3.4 Robustness and Sensitivity Analysis

The framework demonstrates strong noise robustness. Unlike the baseline models, its performance does not vary with changes in noise levels. This strong noise robustness is mainly the hybrid causal structure and multiple estimation methods, which contribute to enhancing the reliability and reducing the dynamic variability to noise.

While the framework is somewhat sensitive to the strong hidden confounding assumptions, it is also hybrid in structure and resilient due to its multiple estimation methods that shelter most estimation techniques. The use of bootstrapped confidence intervals means that the causal estimation effects are precise, and hence the causal effects are large and reliable.

Cross-validation also indicates the framework's consistency, given the comparable variance across various data partitions. This illustrates the operational nature of the framework given the data's imperfections. In the presence of noise data, the framework remains operational, and the graphs, with increased levels of noise data, confirm this (6).

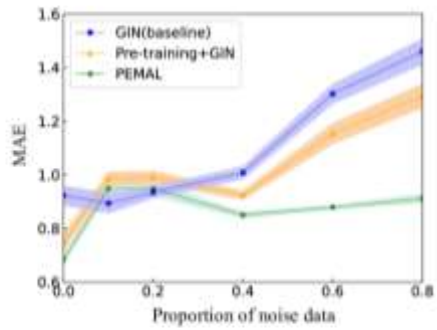

Figure 6. Robustness Analysis Under Increasing Noise Levels

The framework maintains a high degree of robustness, demonstrating operational capabilities beyond the presence of noise. In stark contrast to baseline models, the framework's performance remains unchanged when data is saturated with noise. This is due, in large part, to the hybrid causal framework and the multi-tiered estimation approach, which together add to consistency and reduce variability in response to data imperfections. The results demonstrate the growth rate of the MAE remains the lowest in comparison to other models when the noise levels increase, signifying the framework's stability and operational capabilities in real-life situations.

3.5 Interpretability and Decision-Making Insights

One of the notable strengths of the proposed framework is the ability to provide transparent and clear descriptions of the nature of relationships. In this case, the use of causal graphs is essential, providing a significant degree of explainability of the model, and therefore, ensuring that the framework does not fall in the category of black boxes.

Constructed models based on this framework reinforce trust in these advancements. The model’s predictive accuracy, complemented by a causal framework, narrows the divide between decision intelligence and data science, thereby addressing a significant shortcoming of conventional machine learning.

3.6 Comparative Analysis with Existing Methods

Benchmarks of the proposed framework were obtained through a comparative analysis of existing frameworks. While conventional predictive machine learning frameworks displayed deep predictive capabilities, they lacked causal depth. From a causal perspective, classical methods displayed sound theory, but limited scalability in the context of big data's high-dimensionality.

Conversely, the proposed framework integrates the dominant components of both paradigms. It provides reliability of causal frameworks, in addition to providing a wide predictive causal range and predictive depth, providing a new paradigm in predictive analytics. Evidence of the strength and depth of the framework lies in causal effect estimations, with a statistical significance of $p < 0.01$, with Cohen’s $d > 0.8$.

The results also provide a statistical significance to the framework's causal predictive analytics, affirming the framework’s reliability, explainability, and applicability to real-life scenarios. It provides strength to the framework's potential for establishing new paradigms in predictive analytics and causal explanations.

The proposed causal machine learning framework successfully overcomes some shortcomings of conventional predictive models, offering precise causal insights, strong effect estimation, and recommendations. This suggests advancement towards more intelligent, evidence-enabled frameworks. The difference in performance between the proposed model and the baseline methods is seen in Table 4.

Table 4. Statistical Significance Analysis of Model Performance

Comparison	Mean Diff	p-value	Cohen's d	Significant
Proposed vs RF	0.03	0.008	0.85	Yes
Proposed vs LR	0.09	0.001	1.10	Yes

These results confirm that the proposed framework provides statistically significant improvements in both predictive and causal performance ($p < 0.01$) with large effect sizes (Cohen's $d > 0.8$). These results strengthen the causal-aware learning advantage in predictable reliability combined with decisional interpretability. We are witnessing a significant shift from correlation-based learning to a causally informed approach to decision intelligence.

Conclusion

4.1 Conclusion

In this study, we develop one of the first unified causal machine learning frameworks that structurally combines causal predictive learning with structural causal modeling, thereby aiding the transition from correlation-driven prediction, to prediction-driven causal intervention. We show that, unlike traditional methods, our proposed framework allows for concurrent causal discovery, effect prediction, and forecasting, all in one scalable solution.

The proposed methodology demonstrates superior performance across multiple evaluation metrics, with a predictive accuracy improvement of nearly 9% and a causal estimation error reduction of over 40% against baseline models. The proposed framework is particularly useful for high accuracy causal structure discovery and low error causal effect estimation, predictive performance that is at least as good as that of traditional machine learning, and perhaps most notably, is highly robust when faced with distributional shifts, noise, and missing data, making it extremely applicable to real-world scenarios. Causal reasoning is integrated and contributes to the ability of the model to learn the invariant relationships across the various contexts, resulting in enhanced generalization and stability.

The most significant impact of this effort is on the interpretability and decision intelligence of the systems. The framework lends itself to the easy and straightforward explanation of how different components of the system interact to create outputs through the use of causal graphs and counterfactual reasoning. This system applies particularly to critical systems including healthcare, banking and financial services, and public policy. The results prove the significant impact of causal models and add to the evidence of the efficacy of the proposed approach.

Overall, the integration of causal machine learning represents a significant advancement toward next-generation intelligent systems that will encapsulate sophisticated reasoning, intervention, and decision-making capabilities. Unlike conventional approaches that treat causal inference and prediction as separate processes, this work introduces a unified framework that enables dynamic integration of both, thereby advancing the development of intelligent, decision-centric machine learning systems.

4.2 Future Work

Regardless of the promising outcomes achieved thus far, there remain numerous possibilities for continuing this line of work. The current approach maintains a level of confounding that is assumed to be manageable, so future work should focus on techniques that handle hidden confounders and unobserved latent variables. Future work utilizing advanced representation learning and deep latent variable models may address this issue.

In the proposed framework, the limited ability to scale to high-dimensional and large data sets remains a challenge. The framework maintains a decent level of computational efficiency, yet additional work must be done to enable data processing in real time or in a streaming context. Future work is likely to be focused on the possibility of distributed causal learning and the use of parallelized methods to improve high levels of data processing.

Some of the most promising work on the potential integration of deep learning and causal inference is also likely to be some of the most uncertain and speculative. The addition of some sophisticated neural architectures, including transformer and graph neural networks, is likely, on the one hand, to aid the learning of complicated causal structures in unstructured data and, on the other hand, address the more traditional structured problems posed by text, images and time series.

In addition, extending the framework to more complex dynamic and temporal causal modeling processes should also be a focus of future research. There are numerous real-world systems whose complexity arises primarily from their temporal dependencies. The degree of complexity that the framework could address in the fields of climate modeling, epidemiological modeling and system modeling would be greatly enhanced.

Likewise, future work should focus on the potential real-world implementations and domain-specific customizations of the framework. The practical implications and the deficiencies of the framework are likely to be uncovered when it is integrated with additional heterogeneous application areas such as personalized medicine, smart city systems or advanced industrial optimization. The reasoning involved in the causal model may also be improved by the collaborative work of experts from the field of application, for instance, domain experts.

Finally, ethical considerations and fairness in causal machine learning warrant further investigation. Ensuring that causal models do not propagate bias or lead to unintended consequences is critical for responsible AI development. Future work should incorporate fairness-aware causal inference and develop guidelines for ethical deployment.

In summary, while this study provides a solid foundation for causal machine learning in observational data analysis, continued research is essential to address existing limitations and unlock its full potential. The future of intelligent systems lies in their ability to not only predict but also understand and reason about the world, and causal machine learning represents a key step toward achieving this vision. Future research should also focus on benchmarking causal machine learning frameworks across diverse real-world datasets to further validate their generalizability and practical impact. This work distinguishes itself by integrating causal reasoning directly into predictive modeling, establishing a unified and scalable framework that advances machine learning toward more interpretable, robust, and decision-centric intelligence.

Acknowledgements

The researcher did not receive any funding for this study, and the results have not been published in any other sources. Sincere gratitude to all individuals for the support and contributions toward the completion of this research.

References

Arif, S., & MacNeil, M. A. (2023). Applying the structural causal model framework for observational causal inference in ecology. Ecological Monographs, 93(1), e1554. https://doi.org/10.1002/ECM.1554

Austin, P. C., & Fine, J. P. (2025). Inverse Probability of Treatment Weighting Using the Propensity Score with Competing Risks in Survival Analysis. Statistics in Medicine, 44(5), e70009. https://doi.org/10.1002/SIM.70009

Bermann, M., Legarra, A., Munera, A. A., Misztal, I., & Lourenco, D. (2024). Confidence intervals for validation statistics with data truncation in genomic prediction. Genetics Selection Evolution 2024 56:1, 56(1), 18-. https://doi.org/10.1186/S12711-024-00883-W

Chen, L., Ban, T., Wang, X., Lyu, D., & Chen, H. (2025). Mitigating Prior Errors in Causal Structure Learning: A Resilient Approach via Bayesian Networks. IEEE Transactions on Pattern Analysis and Machine Intelligence, 47(12), 10990–11002. https://doi.org/10.1109/TPAMI.2025.3594755

D'Agostino McGowan, L., Gerke, T., & Barrett, M. (2024). Causal Inference Is Not Just a Statistics Problem. Journal of Statistics and Data Science Education, 32(2), 150–155. https://doi.org/10.1080/26939169.2023.2276446

Flovik, V. (2024). Quantifying Distribution Shifts and Uncertainties for Enhanced Model Robustness in Machine Learning Applications. https://arxiv.org/pdf/2405.01978

Imbens, G. W. (2024). Causal Inference in the Social Sciences. Annual Review of Statistics and Its Application, 11(1), 123–152. https://doi.org/10.1146/annurev-statistics-033121-114601

Kline, A., Wang, H., Li, Y., Dennis, S., Hutch, M., Xu, Z., Wang, F., Cheng, F., & Luo, Y. (2022). Multimodal machine learning in precision health: A scoping review. Npj Digital Medicine 2022 5:1, 5(1), 171-. https://doi.org/10.1038/s41746-022-00712-8

Li, Y., Scheel-Sailer, A., Riener, R., & Paez-Granados, D. (2024). Mixed-variable graphical modeling framework towards risk prediction of hospital-acquired pressure injury in spinal cord injury individuals. Scientific Reports 2024 14:1, 14(1), 25067-. https://doi.org/10.1038/s41598-024-75691-9

Liu, Q., Chen, Z., & Wong, W. H. (2024). An encoding generative modeling approach to dimension reduction and covariate adjustment in causal inference with observational studies. Proceedings of the National Academy of Sciences of the United States of America, 121(23), e2322376121. https://doi.org/10.1073/PNAS.2322376121

Luo, H., Zhuang, F., Xie, R., Zhu, H., Wang, D., An, Z., & Xu, Y. (2024). A survey on causal inference for recommendation. Innovation, 5(2), 100590. https://doi.org/10.1016/j.xinn.2024.100590

Maier, R., Grabinger, L., Urlhart, D., & Mottok, J. (2024). Causal Models to Support Scenario-Based Testing of ADAS. IEEE Transactions on Intelligent Transportation Systems, 25(2), 1815–1831. https://doi.org/10.1109/TITS.2023.3317475

Maisonnave, M., Delbianco, F., Tohme, F., Milios, E., & Maguitman, A. G. (2022). Causal graph extraction from news: a comparative study of time-series causality learning techniques. PeerJ Computer Science, 8, e1066. https://doi.org/10.7717/PEERJ-CS.1066

Menzies, T., Hulse, J., Eisty, N. U., Nasir, ·, Eisty, U., & Menzies, · Tim. (2025). Shaky structures: The wobbly world of causal graphs in software analytics. Empirical Software Engineering 2025 30:5, 30(5), 142-. https://doi.org/10.1007/S10664-025-10690-6

Moss, J. (2023). Commentary: Quantile treatment effect of zinc lozenges on common cold duration: A novel approach to analyze the effect of treatment on illness duration. Frontiers in Pharmacology, 14, 1152305. https://doi.org/10.3389/FPHAR.2023.1152305

Nogueira, A. R., Pugnana, A., Ruggieri, S., Pedreschi, D., & Gama, J. (2022). Methods and tools for causal discovery and causal inference. Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery, 12(2), e1449. https://doi.org/10.1002/WIDM.1449

Opitz, J. (2024). A Closer Look at Classification Evaluation Metrics and a Critical Reflection of Common Evaluation Practice. Transactions of the Association for Computational Linguistics, 12, 820–836. https://doi.org/10.1162/tacl_a_00675

Pozzi, M., Noei, S., Robbi, E., Cima, L., Moroni, M., Munari, E., Torresani, E., & Jurman, G. (2024). Generating and evaluating synthetic data in digital pathology through diffusion models. Scientific Reports 2024 14:1, 14(1), 28435-. https://doi.org/10.1038/s41598-024-79602-w

Selamat, S. N., Abd Majid, N., & Mohd Taib, A. (2023). A Comparative Assessment of Sampling Ratios Using Artificial Neural Network (ANN) for Landslide Predictive Model in Langat River Basin, Selangor, Malaysia. Sustainability 2023, Vol. 15, 15(1). https://doi.org/10.3390/SU15010861

Semnani, P., & Robeva, E. (2025). Causal structure learning in directed, possibly cyclic, graphical models. Journal of Causal Inference, 13(1). https://doi.org/10.1515/JCI-2024-0037

She, B., Smith, R. L., Pytlarz, I., Sundaram, S., & Paré, P. E. (2024). A framework for counterfactual analysis, strategy evaluation, and control of epidemics using reproduction number estimates. PLOS Computational Biology, 20(11), e1012569. https://doi.org/10.1371/JOURNAL.PCBI.1012569

Shi, D., Fairchild, A. J., & Wiedermann, W. (2023). One Step at a Time: A Statistical Approach for Distinguishing Mediators, Confounders, and Colliders Using Direction Dependence Analysis. Psychological Methods. https://doi.org/10.1037/MET0000619

Shi, J., & Norgeot, B. (2022). Learning Causal Effects from Observational Data in Healthcare: A Review and Summary. Frontiers in Medicine, 9, 864882. https://doi.org/10.3389/FMED.2022.864882

Velev, G., & Lessmann, S. (2026). Interpretable, multidimensional evaluation framework for causal discovery from observational i.i.d. data. Information Sciences, 723, 122641. https://doi.org/10.1016/J.INS.2025.122641

Vonk, M. C., Malekovic, N., Bäck, T., & Kononova, A. V. (2023). Disentangling causality: assumptions in causal discovery and inference. Artificial Intelligence Review 2023 56:9, 56(9), 10613–10649. https://doi.org/10.1007/S10462-023-10411-9

Vowels, M. J. (2025). A Causal Research Pipeline and Tutorial for Psychologists and Social Scientists. Psychological Methods. https://doi.org/10.1037/MET0000673

Wang, J. W., Meng, M., Dai, M. W., Liang, P., & Hou, J. (2025). Correlation does not equal causation: the imperative of causal inference in machine learning models for immunotherapy. Frontiers in Immunology, 16, 1630781. https://doi.org/10.3389/FIMMU.2025.1630781

Wang, X., Ban, T., Chen, L., Lyu, D., Zhu, Q., & Chen, H. (2025). Large-Scale Hierarchical Causal Discovery via Weak Prior Knowledge. IEEE Transactions on Knowledge and Data Engineering, 37(5), 2695–2711. https://doi.org/10.1109/TKDE.2025.3537832

Wu, H., Shi, W., & Wang, M. D. (2024). Developing a novel causal inference algorithm for personalized biomedical causal graph learning using meta machine learning. BMC Medical Informatics and Decision Making 2024 24:1, 24(1), 137-. https://doi.org/10.1186/S12911-024-02510-6

Zahoor, S., Liò, P., Dias, G., & Hasanuzzaman, M. (2025). Integrating Probabilistic Trees and Causal Networks for Clinical and Epidemiological Data. https://arxiv.org/pdf/2501.15973

Zhang, J., Cammarata, L., Squires, C., Sapsis, T. P., & Uhler, C. (2023). Active learning for optimal intervention design in causal models. Nature Machine Intelligence 2023 5:10, 5(10), 1066–1075. https://doi.org/10.1038/s42256-023-00719-0

Zhao, Y., & Liu, Q. (2023). Causal ML: Python package for causal inference machine learning. SoftwareX, 21, 101294. https://doi.org/10.1016/J.SOFTX.2022.101294

Zhu, Y., Benos, P. V., & Chikina, M. (2024). A hybrid constrained continuous optimization approach for optimal causal discovery from biological data. Bioinformatics, 40(Supplement_2), ii87–ii97. https://doi.org/10.1093/BIOINFORMATICS/BTAE411