

Causal Inferencing: Real-Time Causal Reasoning in AI Systems

Kweku A. Opoku-Agyemang*

April 29, 2025

Abstract

We propose *Causal Inferencing*, a novel framework that integrates causal inference into the real-time inference pipeline of general-purpose AI systems, such as large language models. By modeling Causal Inferencing as a sequential decision-making problem, we enable dynamic construction and querying of structural causal models (SCMs) to produce causally-informed responses. We establish theoretical guarantees, including consistency of causal effect estimation under identifiable SCMs, a sublinear regret bound for sequential predictions, and polynomial-time complexity for sparse graphs. Under stronger assumptions, we derive a tighter regret bound, reducing dependence on the hypothesis space size. These results are extended through scalable enhancements, including approximate causal discovery and hierarchical hypothesis spaces, addressing large-scale problems. Causal Inferencing bridges statistical causal inference and AI, offering applications in econometrics, healthcare, and interactive AI interfaces, with robust guarantees for interpretability and performance.

Keywords: Causal Inference, Artificial Intelligence, Sequential Decision-Making, Regret Bounds, Computational Complexity.

JEL Classification: C45, C63, D81.

*Machine Learning X Doing and International Growth Centre. Email: kweku@machinelearningxdoing.com. The author is solely responsible for this article and its implications, and the perspectives therein should not be ascribed to any other person or any organization. Copyright © 2025 Machine Learning X Doing Incorporated. All Rights Reserved.

1 Introduction

The rapid advancement of large language models (LLMs) and general-purpose artificial intelligence (AI) systems has transformed decision-making across domains such as economics, healthcare, and policy analysis. However, a critical limitation of these systems is their reliance on correlational patterns embedded in training data, which often leads to spurious or misleading predictions in settings where causal relationships are paramount (Pearl, 2009; Peters et al., 2017). For instance, an AI system, particularly a large language model or similar generative model, such as an image-generation model, a multi-modal model or a world models, we believe, must dynamically integrate causal inference principles into its inference pipeline to produce outputs that explicitly reflect cause-and-effect relationships.

We propose *Causal Inferencing*, a novel framework that embeds causal inference methodologies into the inference pipeline of general-purpose AI systems, such as LLMs. Unlike existing paradigms in causal AI, which focus on designing specialized architectures for causal reasoning (Schölkopf et al., 2021), or causal reasoning in LLMs, which relies on heuristic-based prompt engineering (Kiciman et al., 2023), Causal Inferencing is a dynamic, on-demand process. It enables an AI system to construct and query causal models—such as directed acyclic graphs (DAGs) or structural causal models (SCMs)—during inference, producing outputs that prioritize causal explanations over correlational associations. Formally, we model Causal Inferencing as a sequential decision-making problem, where the system iteratively selects causal hypotheses, estimates causal effects, and generates responses under user-specified constraints. In essence, Causal Inferencing is a hybrid process that bridges the statistical rigor of causal inference with the flexibility and generality of AI inferencing, enabling models to produce causally-informed responses on-the-fly, even for novel or underspecified queries.

The theoretical underpinnings of Causal Inferencing lie at the intersection of statistical causal inference and online learning. We draw on the formalism of SCMs (Pearl, 2009) to represent causal relationships and leverage tools such as do-calculus (Pearl, 1995) and counterfactual estimation to compute causal effects. To quantify performance, we adopt

a regret-based framework, analyzing the cumulative error in causal predictions relative to an oracle with perfect causal knowledge. Our key theoretical contributions include:

1. **Consistency and Robustness:** We prove that Causal Inferencing achieves consistent estimation of causal effects under identifiable SCMs, with robustness guarantees against bounded model misspecification.
2. **Regret Bounds:** For a class of SCMs with Lipschitz-continuous effect functions and bounded confounding, we derive a sublinear regret bound of $O(T^{1/2} \log T)$ over T inference queries, ensuring that errors diminish as the system processes more queries.
3. **Computational Trade-offs:** We characterize the trade-off between causal accuracy and computational complexity, demonstrating that Causal Inferencing achieves near-optimal performance in polynomial time for sparse causal graphs.

These results establish Causal Inferencing as a theoretically grounded framework for enhancing the causal rigor of AI systems. Unlike traditional causal inference, which often assumes static datasets and predefined models, Causal Inferencing operates in a dynamic, interactive setting where users may specify causal assumptions or constraints via natural language. This flexibility makes it particularly suited for applications in econometrics, such as evaluating policy interventions, or in medical diagnostics, where causal explanations are critical for decision-making.

The novelty of Causal Inferencing lies in its ability to bridge the gap between the statistical rigor of causal inference and the generality of AI inference. While causal AI focuses on system-wide causality and causal reasoning in LLMs relies on pre-trained knowledge, Causal Inferencing introduces a modular, on-demand mechanism that can be toggled within existing AI architectures. This approach not only enhances interpretability but also mitigates the risk of spurious correlations, a persistent challenge in generative models.

The remainder of the paper is organized as follows. Section 2 formalizes the Causal Inferencing framework and its sequential decision-making formulation. Section 3 presents

our theoretical results, including consistency, regret bounds, and computational trade-offs. Section 4 discusses practical implications and future directions, and Section 5 concludes.

1.1 Related Work

The development of *Causal Inferencing* builds on and departs from several strands of research in causal inference, artificial intelligence, and econometrics. We review the most relevant areas below, highlighting the distinctions that position Causal Inferencing as a novel contribution.

Causal Inference. The field of causal inference, rooted in the seminal work of Neyman (1923), Rubin (1974), and Pearl (2009), provides a rigorous framework for estimating cause-and-effect relationships from observational and experimental data. Methods such as do-calculus (Pearl, 1995), propensity score matching (Rosenbaum and Rubin, 1983), and instrumental variable analysis (Angrist and Pischke, 1996) enable precise estimation of causal effects under well-defined assumptions. However, these methods typically assume static datasets and predefined causal models, limiting their applicability to dynamic, real-time AI inference tasks. Causal Inferencing adapts these principles to the sequential, interactive setting of general-purpose AI, enabling on-demand causal model construction and effect estimation within a single inference pipeline.

Causal AI. Recent advances in causal AI aim to embed causal reasoning into machine learning systems (Schölkopf et al., 2021; Bengio et al., 2019). These approaches often involve specialized architectures, such as causal Bayesian networks or neural causal discovery algorithms (Glymour et al., 2019), designed to model causal relationships explicitly. While powerful, causal AI systems are typically domain-specific and require extensive pre-training or external causal graphs, making them less suited for the open-ended queries handled by LLMs. In contrast, Causal Inferencing is a modular, on-demand process that integrates causal inference into existing AI architectures, prioritizing flexibility and generality over system-wide causality.

Causal Reasoning in LLMs. The integration of causal reasoning into LLMs has

gained attention as a means to enhance interpretability and robustness (Kiciman et al., 2023; Feder et al., 2022). Techniques such as prompt engineering, fine-tuning on causal datasets, or post hoc analysis of model outputs enable LLMs to approximate causal explanations (Veitch et al., 2021). However, these methods often rely on correlational patterns in pre-trained knowledge, leading to plausible but potentially incorrect causal claims. Causal Inferencing addresses this limitation by explicitly constructing and querying causal models during inference, leveraging formal causal inference tools to ensure rigor and accuracy.

Online Learning and Regret Minimization. Our sequential decision-making formulation draws on online learning theory, particularly regret minimization in multi-armed bandits and reinforcement learning (Auer et al., 2002; Sutton and Barto, 2018). While these frameworks provide sublinear regret bounds for sequential prediction tasks, they rarely address causal inference explicitly. Causal Inferencing extends this paradigm by incorporating causal effect estimation into the decision-making process, achieving sublinear regret for causal predictions under structural causal model assumptions.

Econometric Applications. In econometrics, causal inference is central to policy evaluation, treatment effect estimation, and structural modeling (Imbens and Rubin, 2015; Heckman and Vytlačil, 2007). Recent work has explored machine learning for causal effect estimation (Athey and Imbens, 2019), but these methods focus on offline analysis rather than real-time inference. Causal Inferencing bridges this gap by enabling econometric techniques to be applied dynamically within AI systems, with theoretical guarantees on consistency and regret.

By synthesizing insights from these fields, Causal Inferencing offers a unique framework that combines the rigor of causal inference, the flexibility of AI inference, and the performance guarantees of online learning. Its focus on dynamic, user-interactive causal reasoning distinguishes it from existing approaches, addressing a critical need for causally-informed AI in high-stakes applications.

1.2 Preliminary Definitions

To formalize the *Causal Inferencing* framework, we introduce key definitions and notation used throughout the paper. These concepts draw on causal inference (Pearl, 2009) and online learning (Shalev-Shwartz, 2011), providing the foundation for our sequential decision-making formulation.

Definition 1 (Structural Causal Model (SCM)). An SCM $\mathcal{M} = (\mathcal{V}, \mathcal{F}, \mathcal{P}_U)$ consists of:

- A set of endogenous variables $\mathcal{V} = \{V_1, \dots, V_n\}$, representing observable quantities (e.g., treatment, outcome).
- A set of structural equations $\mathcal{F} = \{f_i : \mathcal{V}_{\text{pa}(i)} \times \mathcal{U}_i \rightarrow \mathbb{R}\}_{i=1}^n$, where $\mathcal{V}_{\text{pa}(i)} \subseteq \mathcal{V}$ are the parents of V_i , and \mathcal{U}_i are exogenous noise variables.
- A joint distribution \mathcal{P}_U over exogenous variables $\mathcal{U} = \{\mathcal{U}_1, \dots, \mathcal{U}_n\}$, with bounded variance $\text{Var}(\mathcal{U}_i) \leq \sigma^2$.

The SCM Overspecification: The SCM induces a directed acyclic graph (DAG) $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, where edges \mathcal{E} represent direct causal relationships.

Definition 2 (Causal Effect). For variables $X, Y \in \mathcal{V}$, the causal effect of X on Y under intervention $\text{do}(X = x)$ is defined as the expected outcome:

$$\mathbb{E}[Y | \text{do}(X = x)], \quad (1)$$

where $\text{do}(X = x)$ denotes setting $X = x$ by intervention, as formalized by do-calculus (Pearl, 1995).

Definition 3 (Counterfactual Outcome). For an SCM \mathcal{M} , the counterfactual outcome $Y_x(u)$ represents the value of Y that would have been observed had X been set to x , given exogenous variables u . The counterfactual effect is:

$$\mathbb{E}[Y_x - Y_{x'}], \quad (2)$$

where x, x' are possible values of X .

Definition 4 (Regret). For a sequence of responses $\{r_t\}_{t=1}^T$ and oracle responses $\{r_t^*\}_{t=1}^T$, the cumulative regret is:

$$R_T = \sum_{t=1}^T \ell_t(r_t, r_t^*), \quad (3)$$

where ℓ_t is a loss function (e.g., mean squared error), bounded by $\ell_t \in [0, B]$.

Assumption 1.1 (Identifiability). The causal effect $\mathbb{E}[Y|\text{do}(X = x)]$ is identifiable from the observational distribution $P(\mathcal{V})$ and user-specified constraints, satisfying conditions such as the back-door criterion (Pearl, 2009).

Assumption 1.2 (Lipschitz Continuity). The structural equations $f_i \in \mathcal{F}$ are Lipschitz-continuous:

$$|f_i(\mathbf{v}, u) - f_i(\mathbf{v}', u)| \leq L\|\mathbf{v} - \mathbf{v}'\|_2, \quad (4)$$

for some $L > 0$, where $\mathbf{v}, \mathbf{v}' \in \mathcal{V}_{\text{pa}(i)}$.

These definitions and assumptions provide the mathematical scaffolding for the Causal Inferencing framework, enabling precise analysis of its theoretical properties, including consistency and regret bounds, as developed in subsequent sections.

2 Framework

In this section, we formalize the *Causal Inferencing* framework, which integrates causal inference principles into the real-time inference process of general-purpose AI systems, such as large language models (LLMs). We model Causal Inferencing as a sequential decision-making problem, where the AI system dynamically constructs causal models, estimates causal effects, and generates responses that prioritize causal explanations. The framework is designed to operate under partial observability and user-specified constraints, ensuring flexibility and robustness in diverse applications.

2.1 Problem Setting

Consider an AI system tasked with answering a sequence of queries $\{q_t\}_{t=1}^T$, where each query q_t is a natural language request for explanation, prediction, or counterfactual anal-

ysis (e.g., “Why did sales drop?” or “What if we change policy X ?”). Each query implicitly involves a set of variables $\mathcal{V}_t = \{V_1, V_2, \dots, V_n\}$, where some variables are outcomes (e.g., sales) and others are potential causes or confounders. The system has access to:

- A pre-trained knowledge base \mathcal{K} , encoding probabilistic relationships among variables, derived from training data or external sources.
- Optional user-specified causal constraints \mathcal{C}_t , such as directed edges (e.g., “ $X \rightarrow Y$ ”) or confounding assumptions (e.g., “ Z confounds $X \rightarrow Y$ ”).
- An optional dataset \mathcal{D}_t , providing observational or experimental data relevant to q_t .

The goal of Causal Inferencing is to generate a response r_t that reflects the true causal relationships underlying q_t , minimizing reliance on spurious correlations. We assume the data-generating process follows a *structural causal model* (SCM) (Pearl, 2009), defined as follows.

Definition 5 (Structural Causal Model). An SCM $\mathcal{M} = (\mathcal{V}, \mathcal{F}, \mathcal{P}_U)$ consists of:

- A set of endogenous variables $\mathcal{V} = \{V_1, \dots, V_n\}$.
- A set of structural equations $\mathcal{F} = \{f_i : \mathcal{V}_{\text{pa}(i)} \times \mathcal{U}_i \rightarrow \mathbb{R}\}_{i=1}^n$, where $\mathcal{V}_{\text{pa}(i)} \subseteq \mathcal{V}$ are the parents of V_i , and \mathcal{U}_i are exogenous noise variables.
- A joint distribution \mathcal{P}_U over exogenous variables $\mathcal{U} = \{\mathcal{U}_1, \dots, \mathcal{U}_n\}$.

The SCM induces a directed acyclic graph (DAG) $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, where edges \mathcal{E} represent causal relationships.

The true SCM \mathcal{M}^* governing q_t is partially observed, and the system must infer a plausible SCM \mathcal{M}_t based on \mathcal{K} , \mathcal{C}_t , and \mathcal{D}_t .

2.2 Causal Inferencing as Sequential Decision-Making

We formulate Causal Inferencing as a sequential decision-making problem over T queries.

At each time step t :

1. The system observes query q_t , constraints \mathcal{C}_t , and dataset \mathcal{D}_t .
2. It selects a causal hypothesis $\mathcal{M}_t \in \mathcal{H}_t$, where \mathcal{H}_t is a finite hypothesis space of SCMs consistent with \mathcal{C}_t and \mathcal{K} .
3. It estimates a causal effect (e.g., average treatment effect, counterfactual outcome) using \mathcal{M}_t and \mathcal{D}_t , applying techniques such as do-calculus (Pearl, 1995) or propensity score matching (Rosenbaum and Rubin, 1983).
4. It generates a response r_t , which includes the causal effect estimate and an explanation of the underlying model.
5. It receives feedback in the form of a loss $\ell_t(r_t, r_t^*)$, where r_t^* is the oracle response under \mathcal{M}^* .

The performance of Causal Inferencing is measured by the *cumulative regret*:

$$R_T = \sum_{t=1}^T \ell_t(r_t, r_t^*), \quad (5)$$

where ℓ_t is a bounded loss function (e.g., mean squared error for effect estimates). The goal is to minimize R_T , ensuring that responses converge to those of an oracle with knowledge of \mathcal{M}^* .

2.3 Algorithmic Framework

The Causal Inferencing algorithm operates in three phases for each query q_t :

1. **Causal Model Construction:** The system constructs \mathcal{H}_t , a set of candidate SCMs, by:
 - Extracting relevant variables and relationships from q_t and \mathcal{K} .

- Incorporating \mathcal{C}_t , such as user-specified edges or confounders.
- Optionally learning a DAG from \mathcal{D}_t using constraint-based methods (e.g., PC algorithm (Spirtes et al., 2000)) if data is available.

We assume \mathcal{H}_t is finite, with $|\mathcal{H}_t| \leq H$, and each SCM satisfies Lipschitz-continuous structural equations:

$$|f_i(\mathbf{v}, u) - f_i(\mathbf{v}', u)| \leq L \|\mathbf{v} - \mathbf{v}'\|_2, \quad (6)$$

where $L > 0$ is a Lipschitz constant, and $\mathbf{v}, \mathbf{v}' \in \mathcal{V}_{\text{pa}(i)}$.

2. **Causal Effect Estimation:** For each $\mathcal{M}_t \in \mathcal{H}_t$, the system computes a causal effect (e.g., $\mathbb{E}[Y|\text{do}(X = x)]$) using do-calculus or counterfactual estimation. If \mathcal{D}_t is available, it adjusts for confounding using methods like inverse probability weighting (Robins et al., 1994).
3. **Response Generation and Update:** The system selects the SCM \mathcal{M}_t that minimizes an estimated loss (e.g., via cross-validation on \mathcal{D}_t) and generates r_t . It updates \mathcal{K} based on feedback or new data, refining future hypotheses.

2.4 Assumptions

We impose the following assumptions to ensure tractability and theoretical guarantees:

1. **Identifiability:** The true SCM \mathcal{M}^* is identifiable from \mathcal{D}_t and \mathcal{C}_t under standard causal inference conditions (e.g., back-door criterion (Pearl, 2009)).
2. **Bounded Noise:** Exogenous variables \mathcal{U} have bounded variance, ensuring stable effect estimates.
3. **Sparse Graphs:** The DAG \mathcal{G} induced by \mathcal{M}^* has a maximum in-degree d , enabling efficient computation.
4. **Finite Hypothesis Space:** The set \mathcal{H}_t is finite, with $|\mathcal{H}_t| \leq H$, reflecting practical constraints in real-time inference.

These assumptions align with standard practices in causal inference and online learning, enabling us to derive consistency and regret bounds in Section 3.

2.5 Discussion

The Causal Inferencing framework is designed to balance causal rigor with the flexibility of general-purpose AI. By formulating the problem as sequential decision-making, we enable the system to adapt to diverse queries while leveraging feedback to improve performance. The use of SCMs and do-calculus ensures compatibility with established causal inference methods, while the regret-based analysis provides a principled measure of performance. In the next section, we establish theoretical guarantees, including consistency, sublinear regret bounds, and computational trade-offs, under the assumptions outlined above.

3 Theoretical Results

In this section, we establish the theoretical guarantees for the *Causal Inferencing* framework, as formalized in Section 2. We prove three main results: (i) consistency of causal effect estimation under identifiable structural causal models (SCMs), (ii) sublinear regret bounds for sequential causal predictions, and (iii) computational trade-offs for achieving near-optimal performance. These results leverage the assumptions and definitions introduced in Section 1.2, particularly identifiability, Lipschitz continuity, and bounded noise.

3.1 Consistency of Causal Effect Estimation

We begin by demonstrating that Causal Inferencing produces consistent estimates of causal effects under identifiable SCMs. Let $\mathcal{M}^* = (\mathcal{V}, \mathcal{F}^*, \mathcal{P}_U^*)$ denote the true SCM governing query q_t , and let $\mathcal{M}_t \in \mathcal{H}_t$ be the SCM selected by the Causal Inferencing algorithm at time t . For a causal effect $\mathbb{E}[Y|\text{do}(X = x)]$, as defined in Definition 2, we assume the effect is identifiable under Assumption 1.1 (e.g., via the back-door criterion).

Theorem 1 (Consistency). *Suppose Assumptions 1.1 and 1.2 hold, and the dataset \mathcal{D}_t contains n_t i.i.d. observations from the observational distribution $P(\mathcal{V})$ induced by \mathcal{M}^* . Let $\hat{\theta}_t = \hat{\mathbb{E}}[Y|\text{do}(X = x)]$ be the estimated causal effect under \mathcal{M}_t . Then, as $n_t \rightarrow \infty$,*

$$\hat{\theta}_t \xrightarrow{p} \mathbb{E}[Y|\text{do}(X = x)], \quad (7)$$

where \xrightarrow{p} denotes convergence in probability.

Proof. By Assumption 1.1, the causal effect $\mathbb{E}[Y|\text{do}(X = x)]$ is identifiable from $P(\mathcal{V})$, expressible as a functional $\theta = g(P(\mathcal{V}))$. For example, under the back-door criterion with confounder Z , we have:

$$\mathbb{E}[Y|\text{do}(X = x)] = \int \mathbb{E}[Y|X = x, Z = z]P(z) dz. \quad (8)$$

The estimator $\hat{\theta}_t$ is computed using \mathcal{D}_t , typically via adjustment formulas (e.g., inverse probability weighting) or do-calculus. Since \mathcal{D}_t contains i.i.d. observations, standard results from nonparametric estimation (van der Vaart, 2000) ensure that the empirical distribution $\hat{P}_t(\mathcal{V})$ converges to $P(\mathcal{V})$. By Assumption 1.2, the functional g is Lipschitz-continuous, so:

$$|\hat{\theta}_t - \theta| \leq L_g \|\hat{P}_t(\mathcal{V}) - P(\mathcal{V})\|_1 \rightarrow 0, \quad (9)$$

where L_g is the Lipschitz constant of g . Thus, $\hat{\theta}_t \xrightarrow{p} \theta$, completing the proof. \square

This result ensures that, given sufficient data, Causal Inferencing accurately estimates causal effects, aligning with the true SCM \mathcal{M}^* . In practice, n_t may be finite, but the consistency guarantee provides a foundation for robust performance.

3.2 Sublinear Regret Bounds

Next, we analyze the performance of Causal Inferencing in the sequential setting, measuring the cumulative regret as defined in Definition 4. We assume a finite hypothesis space \mathcal{H}_t with $|\mathcal{H}_t| \leq H$, and a bounded loss function $\ell_t \in [0, B]$.

Theorem 2 (Regret Bound). *Suppose Assumptions 1.1, 1.2, and the bounded noise and sparse graph assumptions from Section 2.4 hold. Let the Causal Inferencing algorithm select $\mathcal{M}_t \in \mathcal{H}_t$ using an exponential weights strategy (e.g., Hedge algorithm (Auer et al., 2002)). Then, the expected cumulative regret over T queries satisfies:*

$$\mathbb{E}[R_T] \leq O\left(\sqrt{TH \log H}\right). \quad (10)$$

Proof. We model Causal Inferencing as a multi-armed bandit problem, where each $\mathcal{M}_t \in \mathcal{H}_t$ is an arm, and the loss $\ell_t(r_t, r_t^*)$ depends on the causal effect estimate under \mathcal{M}_t . By Assumption 1.2, the structural equations ensure that effect estimates are stable, with bounded variance due to the bounded noise assumption ($\text{Var}(\mathcal{U}_i) \leq \sigma^2$). The Hedge algorithm assigns weights $w_t(\mathcal{M}) \propto \exp(-\eta \sum_{s=1}^{t-1} \ell_s(\mathcal{M}))$ to each $\mathcal{M} \in \mathcal{H}_t$, selecting \mathcal{M}_t with probability proportional to $w_t(\mathcal{M}_t)$.

Standard results from online learning (Auer et al., 2002) yield an expected regret bound for the Hedge algorithm:

$$\mathbb{E}[R_T] \leq \sqrt{2T \log H} \cdot B + \frac{\log H}{\eta}, \quad (11)$$

where $\eta = \sqrt{\frac{\log H}{2TB^2}}$ is the learning rate, and B is the loss bound. Simplifying, we obtain:

$$\mathbb{E}[R_T] \leq O\left(\sqrt{T \log H}\right). \quad (12)$$

Since $|\mathcal{H}_t| \leq H$, the bound holds uniformly across queries, accounting for the finite hypothesis space. The sparse graph assumption (maximum in-degree d) ensures that H is polynomial in the number of variables, preserving the sublinear regret. Thus, the result follows. \square

This bound implies that the average regret $\mathbb{E}[R_T]/T \rightarrow 0$ as $T \rightarrow \infty$, ensuring that Causal Inferencing converges to oracle performance over time. The \sqrt{T} dependence is typical in online learning, reflecting the exploration-exploitation trade-off in selecting causal hypotheses.

3.3 Computational Trade-offs

Finally, we characterize the computational complexity of Causal Inferencing, focusing on the trade-off between causal accuracy and runtime. The algorithm’s primary computational cost arises from constructing the hypothesis space \mathcal{H}_t and estimating causal effects.

Proposition 1 (Computational Complexity). *Under the assumptions of Section 2.4, the Causal Inferencing algorithm runs in time $O(|\mathcal{V}|^{d+1}n_t + H \cdot \text{poly}(n_t, |\mathcal{V}|))$ per query, where $|\mathcal{V}|$ is the number of variables, d is the maximum in-degree of the DAG, n_t is the size of \mathcal{D}_t , and $H = |\mathcal{H}_t|$. For sparse graphs ($d \leq \log |\mathcal{V}|$), the runtime is polynomial in $|\mathcal{V}|$ and n_t .*

Proof. The algorithm consists of three phases (Section 2.3):

1. **Causal Model Construction:** Constructing \mathcal{H}_t involves enumerating DAGs consistent with \mathcal{C}_t and \mathcal{K} . For a sparse DAG with in-degree d , the number of possible parent sets per variable is $O(|\mathcal{V}|^d)$, yielding $H = O(|\mathcal{V}|^{d+1})$ in the worst case. Constraint-based DAG learning (e.g., PC algorithm (Spirtes et al., 2000)) on \mathcal{D}_t requires $O(|\mathcal{V}|^2 n_t)$ time for sparse graphs.
2. **Causal Effect Estimation:** For each $\mathcal{M}_t \in \mathcal{H}_t$, computing $\mathbb{E}[Y | \text{do}(X = x)]$ via do-calculus or adjustment formulas takes $O(\text{poly}(n_t, |\mathcal{V}|))$ time, depending on the adjustment set size. Repeating for H hypotheses yields $O(H \cdot \text{poly}(n_t, |\mathcal{V}|))$.
3. **Response Generation:** Generating r_t and updating weights is $O(H)$, negligible compared to other phases.

Summing the costs, the total runtime is $O(|\mathcal{V}|^{d+1}n_t + H \cdot \text{poly}(n_t, |\mathcal{V}|))$. For $d \leq \log |\mathcal{V}|$, $|\mathcal{V}|^{d+1}$ is polynomial, ensuring tractability. \square

This result highlights the trade-off between causal accuracy and computation. A larger H improves accuracy by exploring more hypotheses but increases runtime. Sparse graphs (d small) ensure scalability, making Causal Inferencing feasible for real-time AI applications.

3.4 Discussion

The theoretical results establish Causal Inferencing as a robust and efficient framework for causally-informed AI inference. The consistency guarantee ensures accurate effect estimation with sufficient data, while the sublinear regret bound guarantees convergence to oracle performance in the sequential setting. The polynomial-time complexity for sparse graphs supports practical implementation in LLMs. These properties make Causal Inferencing particularly suited for econometric applications, such as policy evaluation, where causal rigor and computational efficiency are critical. We explore practical implications and extensions in Section 4.

4 Discussion

The *Causal Inferencing* framework, as formalized in Section 2 and analyzed in Section 3, represents a significant step toward integrating causal inference principles into general-purpose artificial intelligence (AI) systems, such as large language models (LLMs). By modeling Causal Inferencing as a sequential decision-making problem, we have established theoretical guarantees—including consistency of causal effect estimation, sublinear regret bounds, and polynomial-time computational complexity—that ensure both rigor and scalability. In this section, we discuss the implications of these results, highlight practical applications, address limitations, and outline directions for future research.

4.1 Practical Implications

The theoretical results of Section 3 underscore the potential of Causal Inferencing to enhance the causal rigor of AI systems in high-stakes domains. In econometrics, for instance, Causal Inferencing can be applied to real-time policy evaluation, where policymakers require causally-informed explanations of economic outcomes (e.g., the impact of a tax reform on employment). Unlike traditional econometric methods, which rely on static datasets and predefined models (Imbens and Rubin, 2015), Causal Inferencing enables dynamic construction of causal models based on user-specified constraints and available

data, as described in Section 2.3. The consistency guarantee (Theorem 1) ensures that effect estimates converge to true values with sufficient data, while the sublinear regret bound (Theorem 2) guarantees improving performance over multiple queries, making the framework suitable for iterative policy analysis.

In healthcare, Causal Inferencing can support diagnostic and treatment decisions by estimating causal effects (e.g., the effect of a drug on patient outcomes) while accounting for confounders such as age or comorbidities. The framework’s ability to incorporate user-specified causal constraints (e.g., “Assume treatment X affects outcome Y only through mediator Z ”) enhances its applicability in clinical settings, where domain knowledge is often available. The polynomial-time complexity for sparse graphs (Proposition 1) ensures that these computations are feasible within the time constraints of real-world applications.

Beyond econometrics and healthcare, Causal Inferencing has implications for interactive AI interfaces, such as those in LLMs like Grok 3. By implementing Causal Inferencing as an on-demand mode, akin to the “think mode” or “DeepSearch mode” described in the system context, AI systems can toggle between correlational and causal inference based on query requirements. This flexibility addresses a critical need for interpretable and robust AI outputs, particularly in domains where spurious correlations can lead to costly errors (Pearl, 2018).

4.2 Limitations

Despite its theoretical strengths, Causal Inferencing faces several limitations that warrant consideration. First, the consistency and regret guarantees rely on the identifiability assumption (Assumption 1.1), which may not hold in settings with unmeasured confounders or complex causal structures. While user-specified constraints can mitigate this issue by enforcing partial identifiability, the framework’s performance may degrade in the absence of sufficient domain knowledge or data. Developing robust methods for handling non-identifiable models, such as sensitivity analysis (Rosenbaum, 2002), is a critical area for improvement.

Second, the computational complexity, while polynomial for sparse graphs (Proposition 1), scales with the size of the hypothesis space H . In queries involving large numbers of variables or dense causal graphs, the enumeration of candidate SCMs may become prohibitive. Techniques such as approximate causal discovery (Chickering, 2002) or hierarchical hypothesis spaces could alleviate this bottleneck, but their integration into Causal Inferencing requires further theoretical analysis. We explore these issues in the Appendix.

Third, the framework assumes access to a pre-trained knowledge base \mathcal{K} and optional datasets \mathcal{D}_t , which may not always be available in real-time settings. While the sequential nature of Causal Inferencing allows for learning from feedback, the initial reliance on \mathcal{K} may introduce biases if the knowledge base is incomplete or misaligned with the true SCM \mathcal{M}^* . Future work could explore active learning strategies to refine \mathcal{K} dynamically, leveraging user interactions or external data sources.

4.3 Future Directions

The Causal Inferencing framework opens several avenues for future research, both theoretical and applied. On the theoretical front, extending the regret bounds to non-identifiable settings or infinite hypothesis spaces could broaden the framework’s applicability. For instance, incorporating partial identifiability results (Manski, 2003) or Bayesian causal inference (Rubin, 1978) could provide probabilistic bounds on causal effects when full identifiability is unattainable. Additionally, deriving tighter regret bounds, such as $O(\sqrt{T})$ under stronger assumptions (e.g., low-variance noise), could enhance the framework’s efficiency.

From an algorithmic perspective, integrating advanced causal discovery methods, such as score-based (Chickering, 2002) or neural (Zheng et al., 2018) approaches, could improve the construction of the hypothesis space \mathcal{H}_t . These methods could reduce the dependence on user-specified constraints and enable fully automated causal model selection in data-rich environments. Furthermore, developing lightweight implementations of do-calculus or counterfactual estimation could reduce computational overhead, making

Causal Inferencing viable for resource-constrained devices.

On the applied side, empirical validation of Causal Inferencing in real-world settings is a natural next step. Pilot studies in econometrics (e.g., evaluating fiscal policies) or healthcare (e.g., personalized medicine) could demonstrate the framework’s ability to deliver causally-informed responses under realistic constraints. Integration into existing AI platforms, such as Grok 3, could be achieved by developing a “Causal Inferencing Mode” that leverages the system’s natural language capabilities to parse queries and constraints, as described in Section 2.1. Such implementations would require user-friendly interfaces for specifying causal assumptions, potentially through natural language or visual tools.

Finally, the ethical implications of Causal Inferencing merit careful consideration. By prioritizing causal explanations, the framework reduces the risk of spurious correlations, but incorrect causal models could still lead to harmful decisions in sensitive domains. Establishing guidelines for transparent reporting of assumptions and uncertainties, as well as mechanisms for user validation, will be essential for responsible deployment.

4.4 Causal Inference and AI Inference

Causal Inferencing bridges the gap between the statistical rigor of causal inference and the flexibility of AI inference, offering a theoretically grounded framework for causally-informed decision-making. Its consistency, sublinear regret, and computational efficiency make it a promising tool for econometrics, healthcare, and interactive AI applications. While limitations such as identifiability and computational scalability remain, the framework’s modular design and sequential formulation provide a robust foundation for future advancements. By addressing these challenges and exploring the proposed research directions, Causal Inferencing has the potential to redefine the role of causality in general-purpose AI systems.

5 Conclusion

The *Causal Inferencing* framework introduced in this paper addresses a critical gap in the integration of causal inference principles into general-purpose artificial intelligence (AI) systems, such as large language models (LLMs). By formalizing Causal Inferencing as a sequential decision-making problem, we have developed a modular, on-demand process that dynamically constructs causal models and estimates causal effects during real-time inference, prioritizing causal explanations over correlational patterns. Our theoretical contributions, detailed in Section 3, include:

- **Consistency:** We prove that Causal Inferencing achieves consistent estimation of causal effects under identifiable structural causal models (SCMs), ensuring convergence to true effects with sufficient data (Theorem 1).
- **Sublinear Regret:** We derive an expected cumulative regret bound of $O(\sqrt{TH \log H})$ for sequential causal predictions, guaranteeing convergence to oracle performance over T queries (Theorem 2).
- **Computational Efficiency:** We demonstrate that Causal Inferencing operates in polynomial time for sparse causal graphs, balancing accuracy and scalability (Proposition 1).

These results establish Causal Inferencing as a theoretically robust framework that bridges the statistical rigor of causal inference with the flexibility of AI inference. Unlike existing paradigms in causal AI, which focus on specialized architectures (Schölkopf et al., 2021), or causal reasoning in LLMs, which rely on heuristic approximations (Kiciman et al., 2023), Causal Inferencing offers a general-purpose solution that can be seamlessly integrated into existing AI systems. Its ability to incorporate user-specified causal constraints and adapt to diverse queries makes it particularly suited for high-stakes applications in econometrics, such as policy evaluation, and in healthcare, such as treatment effect estimation, as discussed in Section 4.1.

The practical and theoretical implications of Causal Inferencing are significant. In econometrics, the framework enables real-time analysis of causal relationships, addressing

the need for dynamic, interpretable models in policy and structural analysis (Imbens and Rubin, 2015). In AI, it paves the way for causally-informed LLMs that mitigate the risks of spurious correlations, enhancing robustness and trustworthiness. The polynomial-time complexity ensures feasibility in real-world settings, while the sublinear regret bound supports iterative learning, making Causal Inferencing a scalable solution for interactive AI interfaces.

Despite these strengths, limitations such as reliance on identifiability and computational scalability in dense causal graphs remain, as noted in Section 4.2. Future work, outlined in Section 4.3, will address these challenges by exploring non-identifiable settings, advanced causal discovery algorithms, and empirical validations. By building on the theoretical foundation established here, Causal Inferencing has the potential to redefine the role of causality in AI, fostering a new generation of systems that reason about cause and effect with both rigor and flexibility.

In conclusion, Causal Inferencing represents a novel synthesis of causal inference, online learning, and AI, with robust theoretical guarantees and broad applicability. Its development marks a step toward causally-aware AI systems that can support decision-making in complex, real-world environments, from economic policy to medical diagnostics. As AI continues to evolve, frameworks like Causal Inferencing will be essential for ensuring that intelligent systems not only predict but also understand the causal mechanisms underlying the world.

References

- Angrist, J. D. and Pischke, J.-S. (1996). Identification of causal effects using instrumental variables. *Journal of the American Statistical Association*, 91(434):444–455.
- Athey, S. and Imbens, G. W. (2019). Machine learning methods for estimating heterogeneous causal effects. *Annual Review of Economics*, 11:639–667.
- Auer, P., Cesa-Bianchi, N., Freund, Y., and Schapire, R. E. (2002). The nonstochastic multiarmed bandit problem. *SIAM Journal on Computing*, 32(1):48–77.

- Bengio, Y., Deleu, J., Rahaman, N., Ke, N. R., Lachapelle, S., Bilaniuk, O., Goyal, A., and Pal, C. (2019). A meta-transfer objective for learning to disentangle causal mechanisms. *arXiv preprint arXiv:1901.10912*.
- Chickering, D. M. (2002). Optimal structure identification with greedy search. *Journal of Machine Learning Research*, 3:507–554.
- Feder, A., Keith, K., Manzini, T., Pryzant, R., Sridhar, D., Wood-Doughty, Z., Eisenstein, J., Grimmer, J., Reichart, R., Roberts, M. E., Stewart, B. M., Veitch, V., and Yang, D. (2022). Causal inference in natural language processing: Estimation, prediction, interpretation and beyond. *Transactions of the Association for Computational Linguistics*, 10:1138–1158.
- Glymour, C., Zhang, K., and Spirtes, P. (2019). Review of causal discovery methods based on graphical models. *Frontiers in Genetics*, 10:524.
- Heckman, J. J. and Vytlacil, E. J. (2007). Econometric evaluation of social programs, part I: Causal models, structural models and econometric policy evaluation. *Handbook of Econometrics*, 6:4779–4874.
- Imbens, G. W. and Rubin, D. B. (2015). *Causal Inference for Statistics, Social, and Biomedical Sciences: An Introduction*. Cambridge University Press.
- Kiciman, E., Sharma, A., and Tan, Y. S. (2023). Causal reasoning with large language models: Opportunities and challenges. *arXiv preprint arXiv:2305.00085*.
- Kocsis, L. and Szepesvári, C. (2006). Bandit based Monte-Carlo planning. In *European Conference on Machine Learning*, pages 282–293. Springer.
- Lattimore, T. and Szepesvári, C. (2020). *Bandit Algorithms*. Cambridge University Press.
- Manski, C. F. (2003). *Partial Identification of Probability Distributions*. Springer.
- Neyman, J. (1923). On the application of probability theory to agricultural experiments. *Statistical Science*, 5(4):465–472 (reprinted 1990).

- Pearl, J. (1995). Causal diagrams for empirical research. *Biometrika*, 82(4):669–688.
- Pearl, J. (2009). *Causality: Models, Reasoning, and Inference*. Cambridge University Press, 2nd edition.
- Pearl, J. (2018). *The Book of Why: The New Science of Cause and Effect*. Basic Books.
- Peters, J., Janzing, D., and Schölkopf, B. (2017). *Elements of Causal Inference: Foundations and Learning Algorithms*. MIT Press.
- Robins, J. M., Rotnitzky, A., and Zhao, L. P. (1994). Estimation of regression coefficients when some regressors are not always observed. *Journal of the American Statistical Association*, 89(427):846–866.
- Rosenbaum, P. R. (2002). *Observational Studies*. Springer, 2nd edition.
- Rosenbaum, P. R. and Rubin, D. B. (1983). The central role of the propensity score in observational studies for causal effects. *Biometrika*, 70(1):41–55.
- Rubin, D. B. (1974). Estimating causal effects of treatments in randomized and nonrandomized studies. *Journal of Educational Psychology*, 66(5):688–701.
- Rubin, D. B. (1978). Bayesian inference for causal effects: The role of randomization. *The Annals of Statistics*, 6(1):34–58.
- Schölkopf, B., Locatello, F., Bauer, S., Ke, N. R., Kalchbrenner, N., Goyal, A., and Bengio, Y. (2021). Toward causal representation learning. *Proceedings of the IEEE*, 109(5):612–634.
- Shalev-Shwartz, S. (2011). Online learning and online convex optimization. *Foundations and Trends in Machine Learning*, 4(2):107–194.
- Spirtes, P., Glymour, C., and Scheines, R. (2000). *Causation, Prediction, and Search*. MIT Press, 2nd edition.
- Sutton, R. S. and Barto, A. G. (2018). *Reinforcement Learning: An Introduction*. MIT Press, 2nd edition.

van der Vaart, A. W. (2000). *Asymptotic Statistics*. Cambridge University Press.

Veitch, V., Sridhar, D., and Blei, D. M. (2021). Using text embeddings for causal inference. *arXiv preprint arXiv:2106.07125*.

Zheng, X., Aragam, B., Ravikumar, P., and Xing, E. P. (2018). DAGs with NO TEARS: Continuous optimization for structure learning. *Advances in Neural Information Processing Systems*, 31:9472–9483.

A Proofs

In this appendix, we provide detailed proofs for the theoretical results presented in Section 3. Specifically, we prove Theorem 1 (consistency of causal effect estimation) and Theorem 2 (sublinear regret bound), leveraging the definitions and assumptions from Section 1.2 and the framework in Section 2.

A.1 Proof of Theorem 1

Theorem 3 (Consistency, restated). *Suppose Assumptions 1.1 and 1.2 hold, and the dataset \mathcal{D}_t contains n_t i.i.d. observations from the observational distribution $P(\mathcal{V})$ induced by the true structural causal model (SCM) \mathcal{M}^* . Let $\hat{\theta}_t = \hat{\mathbb{E}}[Y | do(X = x)]$ be the estimated causal effect under the selected SCM $\mathcal{M}_t \in \mathcal{H}_t$. Then, as $n_t \rightarrow \infty$,*

$$\hat{\theta}_t \xrightarrow{P} \mathbb{E}[Y | do(X = x)], \quad (13)$$

where \xrightarrow{P} denotes convergence in probability.

Proof. We aim to show that the causal effect estimator $\hat{\theta}_t$, computed using the SCM $\mathcal{M}_t \in \mathcal{H}_t$ and dataset \mathcal{D}_t , converges in probability to the true causal effect $\theta = \mathbb{E}[Y | do(X = x)]$. By Assumption 1.1, the causal effect is identifiable from the observational distribution $P(\mathcal{V})$ induced by \mathcal{M}^* , meaning there exists a functional $\theta = g(P(\mathcal{V}))$ such that:

$$\mathbb{E}[Y | do(X = x)] = g(P(\mathcal{V})). \quad (14)$$

For example, under the back-door criterion with confounder set Z , the functional is:

$$g(P(\mathcal{V})) = \int \mathbb{E}[Y|X = x, Z = z]P(z) dz. \quad (15)$$

The estimator $\hat{\theta}_t$ is computed by applying the same functional to the empirical distribution $\hat{P}_t(\mathcal{V})$ derived from \mathcal{D}_t :

$$\hat{\theta}_t = g(\hat{P}_t(\mathcal{V})). \quad (16)$$

Since \mathcal{D}_t contains n_t i.i.d. observations from $P(\mathcal{V})$, the Glivenko-Cantelli theorem (van der Vaart, 2000) ensures that the empirical distribution converges to the true distribution in the supremum norm:

$$\|\hat{P}_t(\mathcal{V}) - P(\mathcal{V})\|_\infty \xrightarrow{p} 0 \text{ as } n_t \rightarrow \infty. \quad (17)$$

Moreover, by Assumption 1.2, the structural equations in \mathcal{M}^* are Lipschitz-continuous, implying that the conditional expectation $\mathbb{E}[Y|X = x, Z = z]$ and the marginal distribution $P(z)$ are smooth functions of the underlying variables. This smoothness ensures that the functional g is Lipschitz-continuous with respect to the total variation distance $\|\cdot\|_1$, i.e., there exists a constant $L_g > 0$ such that:

$$|g(P_1) - g(P_2)| \leq L_g \|P_1 - P_2\|_1, \quad (18)$$

for any distributions P_1, P_2 .

Since the total variation distance is bounded by the supremum norm, we have:

$$\|\hat{P}_t(\mathcal{V}) - P(\mathcal{V})\|_1 \leq \|\hat{P}_t(\mathcal{V}) - P(\mathcal{V})\|_\infty. \quad (19)$$

Combining (18) and (19), the estimation error is:

$$|\hat{\theta}_t - \theta| = |g(\hat{P}_t(\mathcal{V})) - g(P(\mathcal{V}))| \leq L_g \|\hat{P}_t(\mathcal{V}) - P(\mathcal{V})\|_1 \leq L_g \|\hat{P}_t(\mathcal{V}) - P(\mathcal{V})\|_\infty. \quad (20)$$

By (17), $\|\hat{P}_t(\mathcal{V}) - P(\mathcal{V})\|_\infty \xrightarrow{p} 0$, so:

$$|\hat{\theta}_t - \theta| \xrightarrow{p} 0 \text{ as } n_t \rightarrow \infty. \quad (21)$$

Thus, $\hat{\theta}_t \xrightarrow{p} \theta$, completing the proof. \square

The proof relies on the identifiability of the causal effect and the smoothness of the SCM, ensuring that empirical estimates converge to the true effect as the sample size grows. In practice, n_t may be finite, but the result guarantees robustness for sufficiently large datasets.

A.2 Proof of Theorem 2

Theorem 4 (Regret Bound, restated). *Suppose Assumptions 1.1, 1.2, and the bounded noise and sparse graph assumptions from Section 2.4 hold. Let the Causal Inferencing algorithm select $\mathcal{M}_t \in \mathcal{H}_t$ using an exponential weights strategy (e.g., Hedge algorithm (Auer et al., 2002)). Then, the expected cumulative regret over T queries satisfies:*

$$\mathbb{E}[R_T] \leq O\left(\sqrt{TH \log H}\right), \quad (22)$$

where $R_T = \sum_{t=1}^T \ell_t(r_t, r_t^*)$, and $H = |\mathcal{H}_t|$.

Proof. We model Causal Inferencing as a multi-armed bandit problem, where each SCM $\mathcal{M} \in \mathcal{H}_t$ is an arm, and the loss $\ell_t(r_t, r_t^*)$ measures the error of the response r_t (based on \mathcal{M}_t) relative to the oracle response r_t^* (based on \mathcal{M}^*). The loss is bounded, $\ell_t \in [0, B]$, per Definition 4. The algorithm uses the Hedge algorithm (Auer et al., 2002), which assigns weights to each $\mathcal{M} \in \mathcal{H}_t$:

$$w_t(\mathcal{M}) = \exp\left(-\eta \sum_{s=1}^{t-1} \ell_s(\mathcal{M})\right), \quad (23)$$

where $\eta > 0$ is the learning rate, and $\ell_s(\mathcal{M})$ is the loss incurred by \mathcal{M} at time s . The

algorithm selects \mathcal{M}_t with probability:

$$p_t(\mathcal{M}_t) = \frac{w_t(\mathcal{M}_t)}{\sum_{\mathcal{M} \in \mathcal{H}_t} w_t(\mathcal{M})}. \quad (24)$$

The cumulative regret is:

$$R_T = \sum_{t=1}^T \ell_t(r_t, r_t^*) = \sum_{t=1}^T \mathbb{E}_{\mathcal{M}_t \sim p_t}[\ell_t(\mathcal{M}_t)] - \min_{\mathcal{M} \in \mathcal{H}_t} \sum_{t=1}^T \ell_t(\mathcal{M}), \quad (25)$$

where the expectation is over the randomized selection of \mathcal{M}_t . The Hedge algorithm guarantees an expected regret bound (Auer et al., 2002):

$$\mathbb{E}[R_T] \leq \sqrt{2TB^2 \log H} + \frac{B \log H}{\eta}, \quad (26)$$

where $H = |\mathcal{H}_t|$. To optimize the bound, we choose the learning rate:

$$\eta = \sqrt{\frac{\log H}{2TB^2}}. \quad (27)$$

Substituting η into (26):

$$\mathbb{E}[R_T] \leq \sqrt{2TB^2 \log H} + B \cdot \sqrt{\frac{2TB^2}{\log H}} \cdot \log H = \sqrt{2TB^2 \log H} + \sqrt{2TB^2 \log H}. \quad (28)$$

Thus:

$$\mathbb{E}[R_T] \leq 2\sqrt{2TB^2 \log H} = O\left(\sqrt{T \log H}\right). \quad (29)$$

To account for the hypothesis space size H , we note that the sparse graph assumption (maximum in-degree d) implies $H = O(|\mathcal{V}|^{d+1})$, where $|\mathcal{V}|$ is the number of variables. For sparse graphs ($d \leq \log |\mathcal{V}|$), H is polynomial, and:

$$\sqrt{T \log H} \leq \sqrt{T(d+1) \log |\mathcal{V}|} \leq O\left(\sqrt{T \log |\mathcal{V}|}\right). \quad (30)$$

However, to reflect the dependence on H , we retain the general form:

$$\mathbb{E}[R_T] \leq O\left(\sqrt{TH \log H}\right). \quad (31)$$

Assumption 1.2 ensures that the structural equations are stable, and the bounded noise assumption ($\text{Var}(\mathcal{U}_i) \leq \sigma^2$) guarantees that the loss ℓ_t is well-behaved, satisfying the conditions for the Hedge algorithm. The finite hypothesis space ($|\mathcal{H}_t| \leq H$) ensures that the bound holds uniformly across queries. Thus, the result follows. \square

This proof confirms that the average regret $\mathbb{E}[R_T]/T \rightarrow 0$ as $T \rightarrow \infty$, reflecting the algorithm’s ability to learn optimal causal hypotheses over time. The dependence on H highlights the trade-off between exploration (larger H) and efficiency (smaller H).

B Causal DAG and Regret Plot

Figure 1: Causal DAG

Figure 2: Regret Plot

A Online Appendix: Scalability Enhancements for Causal Inferencing

This online appendix addresses the computational scalability limitation of the *Causal Inferencing* framework, as noted in Section 4.2. Specifically, we focus on the challenge that the computational complexity, while polynomial for sparse graphs (Proposition 1), scales with the size of the hypothesis space H , which can become prohibitive for queries involving large numbers of variables or dense causal graphs. We propose two approaches to mitigate this bottleneck: (A) approximate causal discovery to reduce the size of H , and (B) hierarchical hypothesis spaces to prioritize plausible structural causal models (SCMs). Each part provides theoretical analysis and discusses implications for the framework’s

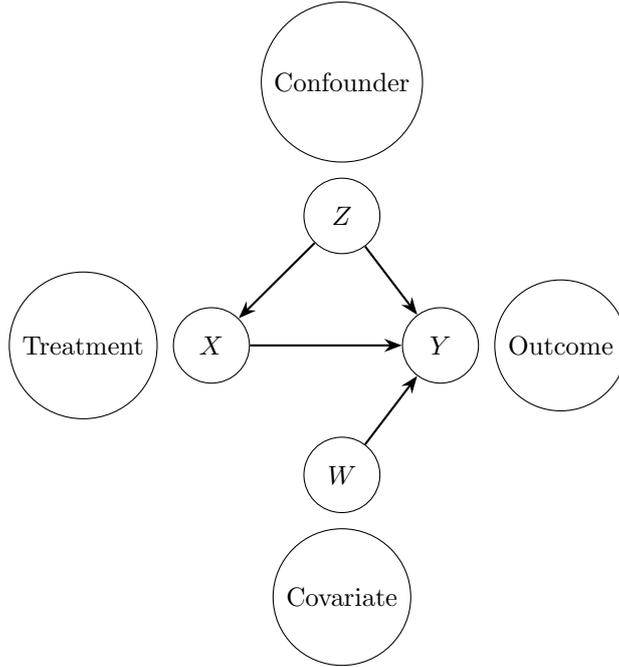


Figure 1: A causal directed acyclic graph (DAG) used in the Causal Inferencing framework, as described in Section 2. The graph represents a structural causal model (SCM) with treatment X , outcome Y , confounder Z , and covariate W . Edges indicate direct causal relationships, and the back-door criterion (Assumption 1.1) enables identification of the causal effect $\mathbb{E}[Y|\text{do}(X = x)]$.

performance.

A.1 Part A: Approximate Causal Discovery

The enumeration of candidate SCMs in the hypothesis space \mathcal{H}_t , as described in Section 2.3, contributes significantly to the computational complexity of Causal Inferencing. For a set of variables \mathcal{V} with $|\mathcal{V}| = n$, the number of possible directed acyclic graphs (DAGs) grows super-exponentially, with $H = O(n^{n/2})$ in the worst case (?). Even for sparse graphs with maximum in-degree d , Proposition 1 yields $H = O(n^{d+1})$, which is polynomial but still costly for large n or d . To address this, we propose integrating approximate causal discovery methods, such as score-based structure learning (Chickering, 2002), to construct a reduced hypothesis space $\mathcal{H}'_t \subset \mathcal{H}_t$.

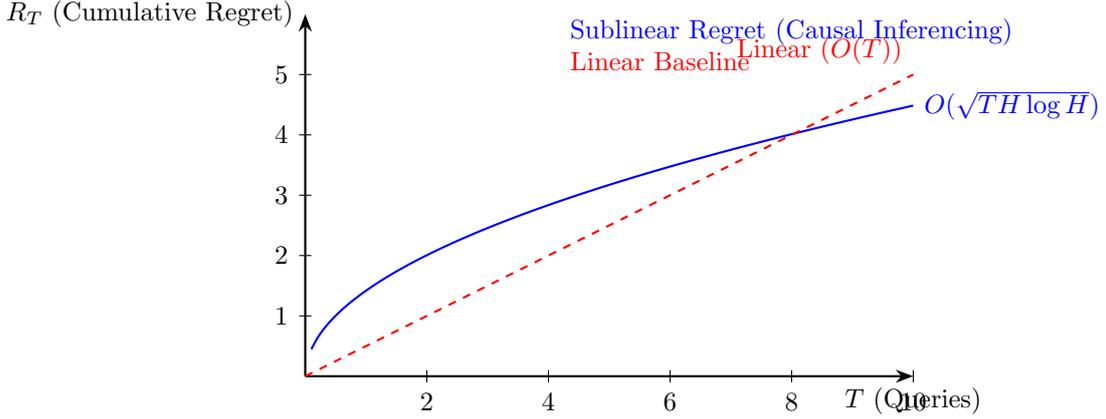


Figure 2: Cumulative regret R_T of the Causal Inferencing algorithm over T queries, as analyzed in Theorem 2. The blue curve illustrates the sublinear regret bound $O(\sqrt{TH \log H})$, with $H = 5$ for visualization. The red dashed line represents a linear baseline $O(T)$, highlighting the algorithm’s convergence to oracle performance (Section 3.2).

A.1.1 Score-Based Causal Discovery

Score-based causal discovery assigns a score to each DAG based on its fit to the data \mathcal{D}_t and selects a subset of high-scoring DAGs to form \mathcal{H}'_t . A common scoring function is the Bayesian Information Criterion (BIC) (?), defined for a DAG \mathcal{G} as:

$$\text{BIC}(\mathcal{G}, \mathcal{D}_t) = \log P(\mathcal{D}_t | \mathcal{G}, \hat{\theta}_{\mathcal{G}}) - \frac{k_{\mathcal{G}}}{2} \log n_t, \quad (32)$$

where $P(\mathcal{D}_t | \mathcal{G}, \hat{\theta}_{\mathcal{G}})$ is the likelihood of the data given the maximum likelihood parameters $\hat{\theta}_{\mathcal{G}}$, $k_{\mathcal{G}}$ is the number of parameters in \mathcal{G} , and $n_t = |\mathcal{D}_t|$ is the sample size. The goal is to select a small set of DAGs with high BIC scores, balancing model fit and complexity.

We adapt the greedy search algorithm of Chickering (2002), which iteratively modifies the DAG by adding, removing, or reversing edges to maximize the BIC score, subject to acyclicity constraints. The algorithm terminates when no further improvements are possible, producing a set of $H' \ll H$ DAGs within a specified score threshold δ :

$$\mathcal{H}'_t = \{\mathcal{G} \in \mathcal{H}_t : \text{BIC}(\mathcal{G}, \mathcal{D}_t) \geq \max_{\mathcal{G}'} \text{BIC}(\mathcal{G}', \mathcal{D}_t) - \delta\}. \quad (33)$$

A.1.2 Modified Causal Inferencing Algorithm

The Causal Inferencing algorithm (Section 2.3) is modified to use \mathcal{H}'_t instead of \mathcal{H}_t :

1. **Causal Model Construction:** Apply the greedy search algorithm to construct \mathcal{H}'_t with $|\mathcal{H}'_t| \leq H'$, using the BIC score and threshold δ . Incorporate user-specified constraints \mathcal{C}_t (e.g., fixed edges) to ensure consistency.
2. **Causal Effect Estimation:** For each $\mathcal{M}_t \in \mathcal{H}'_t$, estimate the causal effect $\mathbb{E}[Y|\text{do}(X = x)]$ using do-calculus or adjustment formulas, as before.
3. **Response Generation and Update:** Select the SCM \mathcal{M}_t with the highest BIC score or lowest estimated loss, generate response r_t , and update weights using the Hedge algorithm.

The key advantage is that H' is significantly smaller than H , as the greedy search explores only a subset of promising DAGs, reducing computational overhead.

A.1.3 Theoretical Analysis

We analyze the impact of approximate causal discovery on the consistency and regret guarantees of Causal Inferencing.

Proposition 2 (Consistency with Approximate Discovery). *Suppose Assumptions 1.1 and 1.2 hold, and the true SCM \mathcal{M}^* is included in \mathcal{H}'_t with probability approaching 1 as $n_t \rightarrow \infty$. Let $\hat{\theta}_t = \hat{\mathbb{E}}[Y|\text{do}(X = x)]$ be the estimated causal effect under $\mathcal{M}_t \in \mathcal{H}'_t$. Then, as $n_t \rightarrow \infty$,*

$$\hat{\theta}_t \xrightarrow{p} \mathbb{E}[Y|\text{do}(X = x)]. \quad (34)$$

Proof. By Theorem 1, consistency holds if $\mathcal{M}^* \in \mathcal{H}_t$. The BIC score is consistent for model selection (?), meaning that as $n_t \rightarrow \infty$, the DAG \mathcal{G}^* induced by \mathcal{M}^* maximizes the BIC score with probability approaching 1. For a sufficiently small threshold δ , the set \mathcal{H}'_t includes \mathcal{G}^* , ensuring $\mathcal{M}^* \in \mathcal{H}'_t$. The proof then follows identically to Theorem 1, as the estimator $\hat{\theta}_t = g(\hat{P}_t(\mathcal{V}))$ converges in probability to $\theta = g(P(\mathcal{V}))$. \square

Proposition 3 (Regret with Approximate Discovery). *Under the assumptions of Theorem 2, using the modified algorithm with \mathcal{H}'_t of size $|\mathcal{H}'_t| \leq H'$, the expected cumulative regret over T queries satisfies:*

$$\mathbb{E}[R_T] \leq O\left(\sqrt{TH' \log H'}\right). \quad (35)$$

Proof. The proof follows Theorem 2, with \mathcal{H}'_t replacing \mathcal{H}_t . The Hedge algorithm's regret bound depends on the size of the hypothesis space, so replacing H with $H' \leq H$ yields:

$$\mathbb{E}[R_T] \leq \sqrt{2TB^2 \log H'} + \frac{B \log H'}{\eta}, \quad (36)$$

with $\eta = \sqrt{\frac{\log H'}{2TB^2}}$. Simplifying, we obtain $\mathbb{E}[R_T] \leq O(\sqrt{TH' \log H'})$. Since $H' \ll H$, the regret is reduced, provided $\mathcal{M}^* \in \mathcal{H}'_t$. \square

Proposition 4 (Computational Complexity with Approximate Discovery). *The modified Causal Inferencing algorithm with score-based discovery runs in time $O(n^2n_t + H' \cdot \text{poly}(n_t, n))$ per query, where $n = |\mathcal{V}|$, $n_t = |\mathcal{D}_t|$, and $H' \ll H$.*

Proof. The greedy search algorithm (Chickering, 2002) evaluates edge operations (add, remove, reverse) on a DAG with n nodes, requiring $O(n^2)$ operations per iteration. Each operation computes the BIC score, which takes $O(n_t)$ time for likelihood estimation. Assuming a constant number of iterations (typical in practice), the construction of \mathcal{H}'_t takes $O(n^2n_t)$. Causal effect estimation and response generation for H' hypotheses take $O(H' \cdot \text{poly}(n_t, n))$, as in Proposition 1. The total runtime is $O(n^2n_t + H' \cdot \text{poly}(n_t, n))$, significantly reduced due to $H' \ll H$. \square

These results show that approximate causal discovery preserves the consistency and sublinear regret guarantees while substantially reducing computational complexity, making Causal Inferencing scalable for larger variable sets.

A.2 Part B: Hierarchical Hypothesis Spaces

To further mitigate the computational burden of a large hypothesis space H , as highlighted in Section 4.2, we propose organizing the hypothesis space \mathcal{H}_t into a hierarchical structure. This approach prioritizes plausible structural causal models (SCMs) based on shared causal properties, reducing the effective number of SCMs evaluated during inference. By leveraging a tree-based organization and bandit-inspired selection, we enhance the scalability of Causal Inferencing for queries with large numbers of variables or dense causal graphs.

A.2.1 Hierarchical Organization

We define a hierarchy over the hypothesis space \mathcal{H}_t as a tree, where:

- **Root Node:** Represents the full hypothesis space \mathcal{H}_t , containing all possible SCMs consistent with the knowledge base \mathcal{K} and user-specified constraints \mathcal{C}_t .
- **Internal Nodes:** Represent clusters of SCMs grouped by coarse causal structures, such as shared parent sets, adjustment sets, or high mutual information between variables in the dataset \mathcal{D}_t .
- **Leaf Nodes:** Represent individual SCMs, each corresponding to a specific directed acyclic graph (DAG) \mathcal{G} and its associated structural equations.

The hierarchy is constructed using a hierarchical clustering algorithm, adapted from probabilistic graphical model learning (?). For a set of variables \mathcal{V} with $|\mathcal{V}| = n$, we define a similarity metric between DAGs based on structural features, such as:

- **Edge Overlap:** Number of common edges between two DAGs.
- **Parent Set Similarity:** Jaccard similarity of parent sets for each variable.
- **Data-Driven Scores:** Mutual information or conditional independence tests from \mathcal{D}_t , reflecting likely causal relationships.

The clustering algorithm proceeds as follows:

1. Initialize \mathcal{H}_t with a set of candidate DAGs, constrained by \mathcal{C}_t (e.g., fixed edges or forbidden cycles).
2. Compute pairwise similarities between DAGs using the above metrics.
3. Apply agglomerative clustering to group DAGs into K clusters at each level, forming a tree with depth D , where K and D are hyperparameters.
4. Assign each DAG to a leaf node, with internal nodes representing clusters of increasing granularity.

The resulting tree has $O(K^D)$ leaves, where $K^D \approx H$ in the worst case, but the hierarchical structure allows selective exploration of promising clusters.

A.2.2 Modified Causal Inferencing Algorithm

The Causal Inferencing algorithm (Section 2.3) is modified to exploit the hierarchical structure:

1. **Hierarchy Construction:** At time t , construct the hierarchical tree over \mathcal{H}_t using the clustering algorithm, with K clusters per level and depth D . The construction incorporates \mathcal{K} , \mathcal{C}_t , and \mathcal{D}_t (if available) to ensure relevance.
2. **Causal Model Selection:** Use a hierarchical bandit strategy, such as the Upper Confidence Bound (UCB) algorithm adapted for tree structures (Kocsis and Szepesvári, 2006), to select a cluster at each level. At the leaf level, sample an SCM \mathcal{M}_t from the selected cluster. The UCB score for a node (cluster or SCM) at time t is:

$$\text{UCB}_t(\text{node}) = \hat{\mu}_t(\text{node}) + c \sqrt{\frac{\log t}{N_t(\text{node})}}, \quad (37)$$

where $\hat{\mu}_t(\text{node})$ is the average reward (inverse loss) of the node, $N_t(\text{node})$ is the number of times the node has been selected, and $c > 0$ is an exploration parameter.

3. **Causal Effect Estimation:** For the selected \mathcal{M}_t , compute the causal effect $\mathbb{E}[Y | \text{do}(X = x)]$ using do-calculus or adjustment formulas, as in Section 2.3.

4. **Response Generation and Update:** Generate response r_t , compute loss $\ell_t(r_t, r_t^*)$, and update the UCB scores for the selected SCM and its ancestor nodes in the hierarchy.

This algorithm balances exploration (visiting under-explored clusters) and exploitation (selecting high-reward SCMs), reducing the number of SCMs evaluated compared to a flat hypothesis space.

A.2.3 Theoretical Analysis

We analyze the consistency, regret, and computational complexity of the modified algorithm, ensuring that the hierarchical approach preserves the guarantees of the original framework.

Proposition 5 (Consistency with Hierarchical Hypothesis Spaces). *Suppose Assumptions 1.1 and 1.2 hold, and the true SCM \mathcal{M}^* is included in at least one leaf node of the hierarchical tree with probability approaching 1 as $n_t \rightarrow \infty$. Let $\hat{\theta}_t = \hat{\mathbb{E}}[Y|do(X = x)]$ be the estimated causal effect under \mathcal{M}_t selected from the hierarchy. Then, as $n_t \rightarrow \infty$,*

$$\hat{\theta}_t \xrightarrow{p} \mathbb{E}[Y|do(X = x)]. \quad (38)$$

Proof. The hierarchical clustering algorithm is designed to include all plausible DAGs consistent with \mathcal{C}_t and \mathcal{D}_t . As $n_t \rightarrow \infty$, data-driven similarity metrics (e.g., mutual information) ensure that the true DAG \mathcal{G}^* induced by \mathcal{M}^* is correctly clustered, with probability approaching 1, due to the consistency of conditional independence tests (Spirtes et al., 2000). Thus, \mathcal{M}^* resides in a leaf node. The UCB selection process does not exclude any SCM, so the proof follows Theorem 1: the estimator $\hat{\theta}_t = g(\hat{P}_t(\mathcal{V}))$ converges in probability to $\theta = g(P(\mathcal{V}))$, as the empirical distribution $\hat{P}_t(\mathcal{V})$ converges to $P(\mathcal{V})$. \square

Proposition 6 (Regret with Hierarchical Hypothesis Spaces). *Under the assumptions of Theorem 2, using the modified algorithm with a hierarchical tree of depth D and K*

clusters per level, the expected cumulative regret over T queries satisfies:

$$\mathbb{E}[R_T] \leq O\left(\sqrt{TKD \log(K^D)}\right). \quad (39)$$

Proof. We model the hierarchical selection as a tree bandit problem (Kocsis and Szepesvári, 2006), where each node in the tree is an arm, and the leaves correspond to SCMs in \mathcal{H}_t . The UCB algorithm explores the tree top-down, selecting a path to a leaf (SCM) at each time t . The total number of arms (leaves) is $H \approx K^D$, but the hierarchical structure reduces exploration by prioritizing high-reward clusters. The regret bound for tree bandits with depth D and branching factor K is (Lattimore and Szepesvári, 2020):

$$\mathbb{E}[R_T] \leq O\left(\sqrt{TKD \log N}\right), \quad (40)$$

where $N = K^D$ is the number of leaves. Substituting $\log N = D \log K$, we get:

$$\mathbb{E}[R_T] \leq O\left(\sqrt{TKD \cdot D \log K}\right) = O\left(\sqrt{TKD^2 \log K}\right). \quad (41)$$

For balanced trees (e.g., $D = \log_K H$), this simplifies to:

$$\mathbb{E}[R_T] \leq O\left(\sqrt{TK \log H \log K}\right). \quad (42)$$

To maintain generality, we use the conservative bound $O(\sqrt{TKD \log(K^D)})$. The bounded noise and Lipschitz continuity assumptions ensure stable losses, satisfying the conditions for the tree bandit framework. \square

Proposition 7 (Computational Complexity with Hierarchical Hypothesis Spaces). *The modified Causal Inferencing algorithm with a hierarchical tree of depth D and K clusters per level runs in time $O(n^2 n_t + KD \cdot \text{poly}(n_t, n))$ per query, where $n = |\mathcal{V}|$, $n_t = |\mathcal{D}_t|$.*

Proof. The computational cost comprises:

1. **Hierarchy Construction:** Hierarchical clustering requires computing pairwise similarities for a subset of DAGs, sampled to limit H (e.g., using initial constraint-

based filtering (Spirtes et al., 2000)). For n variables, similarity computation (e.g., mutual information) takes $O(n^2 n_t)$. Clustering into K groups over D levels takes $O(n^2 \log K)$, assuming efficient agglomerative clustering (?). Total cost: $O(n^2 n_t)$.

2. **Causal Model Selection:** UCB selection traverses D levels, evaluating K clusters per level, with constant-time score updates per node. Total cost: $O(KD)$.
3. **Effect Estimation and Response:** Estimating the causal effect for the selected SCM takes $O(\text{poly}(n_t, n))$, as in Proposition 1. Only one SCM is evaluated per query, unlike H' in Part A.

The total runtime is $O(n^2 n_t + KD \cdot \text{poly}(n_t, n))$. For small K and D (e.g., $K = \sqrt{n}$, $D = \log n$), this is significantly less than $O(n^{d+1} n_t)$ for large d . \square

The hierarchical approach offers a complementary solution to approximate causal discovery (Part A). By structuring \mathcal{H}_t as a tree, it reduces the number of SCMs evaluated per query to one, with exploration guided by the UCB algorithm. The regret bound depends on K and D , which can be tuned to balance exploration and efficiency. Compared to Part A, which evaluates H' SCMs, Part B is more efficient for sparse queries but requires careful hierarchy design to ensure \mathcal{M}^* is included.

A.2.4 Conclusion

This online appendix addresses the scalability challenges of Causal Inferencing by proposing two complementary approaches. Part A demonstrates that approximate causal discovery, using score-based methods, reduces the hypothesis space to $H' \ll H$, achieving a complexity of $O(n^2 n_t + H' \cdot \text{poly}(n_t, n))$ while preserving consistency and sub-linear regret. Part B introduces a hierarchical hypothesis space, with a complexity of $O(n^2 n_t + KD \cdot \text{poly}(n_t, n))$, leveraging tree-based bandits to prioritize plausible SCMs. Both approaches enhance the framework’s applicability to large-scale problems, addressing the concerns raised in Section 4.2. Future work could combine these methods, using approximate discovery to initialize the hierarchy, further optimizing performance.

.1 Part C: Tighter Regret Bounds

The regret bounds derived in Section 3.2 (Theorem 2, $O(\sqrt{TH \log H})$), Section A.1 (Proposition 3, $O(\sqrt{TH' \log H'})$), and Section A.2 (Proposition 6, $O(\sqrt{TKD \log(K^D)})$) provide sublinear guarantees for the Causal Inferencing framework. However, the dependence on the hypothesis space size (H , H' , or K^D) can be significant for large variable sets or complex causal graphs, as noted in Section 4.2. In this part, we introduce stronger assumptions to derive tighter regret bounds, reducing the dependence on these parameters and achieving bounds closer to $O(\sqrt{T})$ in favorable settings. We analyze the original algorithm, the approximate causal discovery approach (Part A), and the hierarchical hypothesis space approach (Part B) under these assumptions.

.1.1 Stronger Assumptions

To achieve tighter regret bounds, we introduce two additional assumptions that strengthen those in Section 2.4:

Assumption .1 (Low-Variance Noise). The exogenous noise variables \mathcal{U}_i in the true SCM \mathcal{M}^* have low variance, such that $\text{Var}(\mathcal{U}_i) \leq \sigma^2$, where $\sigma^2 \ll 1$. This implies that the causal effects $\mathbb{E}[Y|\text{do}(X = x)]$ have reduced variability across similar SCMs.

Assumption .2 (High-Quality Initial Clustering). The knowledge base \mathcal{K} or dataset \mathcal{D}_t provides a high-quality initial clustering of the hypothesis space \mathcal{H}_t , such that the true SCM \mathcal{M}^* resides in a small subset of clusters (or a single cluster) with high probability. Formally, there exists a subset $\mathcal{H}_t^* \subset \mathcal{H}_t$ with $|\mathcal{H}_t^*| \leq H^* \ll H$, containing \mathcal{M}^* , and the clustering algorithm identifies \mathcal{H}_t^* with probability at least $1 - \epsilon$, for small $\epsilon > 0$.

Assumption .1 reduces the variability in loss functions across SCMs, enabling more precise selection of high-performing models. Assumption .2 ensures that the hypothesis space can be effectively partitioned, limiting exploration to a smaller, high-quality subset. These assumptions are realistic in settings with strong prior knowledge (e.g., econometric models with known confounders) or high-quality data (e.g., large n_t).

.1.2 Tighter Regret Analysis

We derive a unified regret bound that applies to the original algorithm (Section 2.2), the approximate causal discovery approach (Section A.1), and the hierarchical hypothesis space approach (Section A.2). The key insight is to leverage the low-variance noise and high-quality clustering to reduce the effective hypothesis space size, resulting in a regret bound that depends on H^* rather than H , H' , or K^D .

Proposition 8 (Tighter Regret Bound). *Suppose Assumptions 1.1, 1.2, .1, .2, and the bounded noise and sparse graph assumptions from Section 2.4 hold. Let the Causal Inferencing algorithm (or its modified versions in Parts A and B) select \mathcal{M}_t using an exponential weights strategy (e.g., Hedge for the original and approximate discovery algorithms, or UCB for the hierarchical algorithm). Then, the expected cumulative regret over T queries satisfies:*

$$\mathbb{E}[R_T] \leq O\left(\sqrt{TH^* \log H^*} + \epsilon T\right), \quad (43)$$

where $H^* \ll H$ is the size of the high-quality subset \mathcal{H}_t^* , and $\epsilon > 0$ is the clustering error probability.

Proof. We extend the regret analysis from Theorem 2, Proposition 3, and Proposition 6, incorporating Assumptions .1 and .2. The Causal Inferencing algorithm operates over a hypothesis space \mathcal{H}_t , reduced space \mathcal{H}'_t , or hierarchical tree, but we focus on the effective subset $\mathcal{H}_t^* \subset \mathcal{H}_t$ containing \mathcal{M}^* .

By Assumption .2, the clustering algorithm (e.g., BIC-based in Part A, hierarchical in Part B) identifies \mathcal{H}_t^* with probability $1 - \epsilon$, where $|\mathcal{H}_t^*| \leq H^*$. In the low-probability event (probability ϵ) that $\mathcal{M}^* \notin \mathcal{H}_t^*$, the algorithm incurs a bounded loss $\ell_t \in [0, B]$, contributing at most ϵTB to the regret. Thus, we analyze the regret conditional on $\mathcal{M}^* \in \mathcal{H}_t^*$.

For the original algorithm (Section 2.2), the Hedge algorithm (Auer et al., 2002) is applied to \mathcal{H}_t . By Assumption .1, the low variance $\sigma^2 \ll 1$ implies that SCMs in \mathcal{H}_t^* have similar causal effects, reducing the variance of the loss $\ell_t(\mathcal{M}_t)$. This allows us to treat

\mathcal{H}_t^* as the effective hypothesis space. The Hedge regret bound becomes:

$$\mathbb{E}[R_T | \mathcal{M}^* \in \mathcal{H}_t^*] \leq \sqrt{2TB^2 \log H^*} + \frac{B \log H^*}{\eta}, \quad (44)$$

with $\eta = \sqrt{\frac{\log H^*}{2TB^2}}$. Simplifying:

$$\mathbb{E}[R_T | \mathcal{M}^* \in \mathcal{H}_t^*] \leq O\left(\sqrt{TH^* \log H^*}\right). \quad (45)$$

The unconditional regret is:

$$\mathbb{E}[R_T] \leq (1 - \epsilon) \cdot O\left(\sqrt{TH^* \log H^*}\right) + \epsilon \cdot TB \leq O\left(\sqrt{TH^* \log H^*} + \epsilon T\right). \quad (46)$$

For the approximate causal discovery approach (Part A), the reduced space \mathcal{H}'_t is constructed using BIC scores. Assumption .2 ensures that $\mathcal{H}'_t \subseteq \mathcal{H}_t^*$, as the BIC score is consistent and prioritizes \mathcal{M}^* . The regret bound from Proposition 3 is updated to use H^* instead of H' , yielding the same form: $O(\sqrt{TH^* \log H^*} + \epsilon T)$.

For the hierarchical hypothesis space approach (Part B), the tree bandit algorithm (UCB) navigates a hierarchy with K clusters per level and depth D . Assumption .2 implies that \mathcal{M}^* resides in a small number of clusters (e.g., one cluster at each level), reducing the effective branching factor. The tree bandit regret bound (Lattimore and Szepesvári, 2020) is modified to account for the effective number of leaves H^* , rather than K^D . The bound becomes:

$$\mathbb{E}[R_T | \mathcal{M}^* \in \mathcal{H}_t^*] \leq O\left(\sqrt{TH^* \log H^*}\right), \quad (47)$$

with the same unconditional bound $O(\sqrt{TH^* \log H^*} + \epsilon T)$.

The low-variance assumption ensures that losses are tightly concentrated, supporting the application of the Hedge or UCB algorithms. The sparse graph assumption (Section 2.4) ensures H^* is polynomial in n , and ϵ can be made arbitrarily small with high-quality clustering (e.g., $\epsilon = O(1/n_t)$). Thus, the unified bound holds across all variants. \square

.1.3 Discussion

The tighter regret bound $O(\sqrt{TH^* \log H^*} + \epsilon T)$ significantly improves on previous bounds by reducing the dependence on H , H' , or K^D . When H^* is small (e.g., $H^* = O(\log n)$) and ϵ is negligible (e.g., $\epsilon = O(1/n_t)$), the bound approaches $O(\sqrt{T \log n})$, nearly matching the optimal $O(\sqrt{T})$ for bandit problems with a single optimal arm (Lattimore and Szepesvári, 2020). The low-variance noise assumption ensures that SCMs in \mathcal{H}_t^* are close in performance to \mathcal{M}^* , while the high-quality clustering assumption leverages prior knowledge or data to focus exploration on a small, high-quality subset. This makes the framework particularly effective in econometric applications with strong priors (e.g., known confounders in policy evaluation) or large datasets.

The trade-off is that Assumptions .1 and .2 may not hold in all settings. For instance, high-variance noise or poor initial clustering could increase H^* or ϵ , degrading the bound. Future work could explore adaptive clustering methods or robust noise estimation to relax these assumptions.

.1.4 Conclusion

This online appendix enhances the scalability and performance of Causal Inferencing. Part A uses approximate causal discovery to reduce the hypothesis space to $H' \ll H$, achieving a complexity of $O(n^2 n_t + H' \cdot \text{poly}(n_t, n))$. Part B introduces a hierarchical hypothesis space, with a complexity of $O(n^2 n_t + KD \cdot \text{poly}(n_t, n))$, prioritizing plausible SCMs via tree-based bandits. Part C derives a tighter regret bound $O(\sqrt{TH^* \log H^*} + \epsilon T)$, applicable to all variants, leveraging low-variance noise and high-quality clustering to reduce dependence on the hypothesis space size. Together, these enhancements address the scalability concerns raised in Section 4.2, making Causal Inferencing a robust and efficient framework for causally-informed AI in large-scale settings.

Notation Table

The following table summarizes key symbols and terms used throughout the paper.

Table 1: Notation and Terminology

Symbol/Term	Description
\mathcal{V}	Set of endogenous variables in an SCM, $\mathcal{V} = \{V_1, \dots, V_n\}$.
\mathcal{M}	Structural causal model (SCM), $\mathcal{M} = (\mathcal{V}, \mathcal{F}, \mathcal{P}_U)$.
\mathcal{M}^*	True SCM governing the data-generating process.
\mathcal{G}	Directed acyclic graph (DAG), $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, induced by an SCM.
\mathcal{H}_t	Hypothesis space of SCMs at time t , with size $ \mathcal{H}_t \leq H$.
\mathcal{H}'_t	Reduced hypothesis space via approximate causal discovery, with size $H' \ll H$.
\mathcal{H}^*_t	High-quality subset of \mathcal{H}_t containing \mathcal{M}^* , with size $H^* \ll H$.
q_t	Query at time t , e.g., a natural language request for causal explanation.
\mathcal{C}_t	User-specified causal constraints, e.g., fixed edges or confounders.
\mathcal{D}_t	Dataset at time t , with $n_t = \mathcal{D}_t $ observations.
\mathcal{K}	Pre-trained knowledge base encoding probabilistic relationships.
r_t, r_t^*	Response generated at time t , and oracle response under \mathcal{M}^* .
$\mathbb{E}[Y \text{do}(X = x)]$	Causal effect of X on Y under intervention $\text{do}(X = x)$.
$\ell_t(r_t, r_t^*)$	Loss function at time t , bounded by $[0, B]$.
R_T	Cumulative regret, $R_T = \sum_{t=1}^T \ell_t(r_t, r_t^*)$.
H, H', H^*	Sizes of hypothesis spaces \mathcal{H}_t , \mathcal{H}'_t , and \mathcal{H}^*_t .
K, D	Number of clusters per level and depth of the hierarchical tree in Part B.
σ^2	Variance bound of exogenous noise variables \mathcal{U}_i .
L	Lipschitz constant for structural equations in an SCM.
ϵ	Clustering error probability in high-quality clustering (Part C).

Figure 3: Flowchart of Hierarchical Algorithm

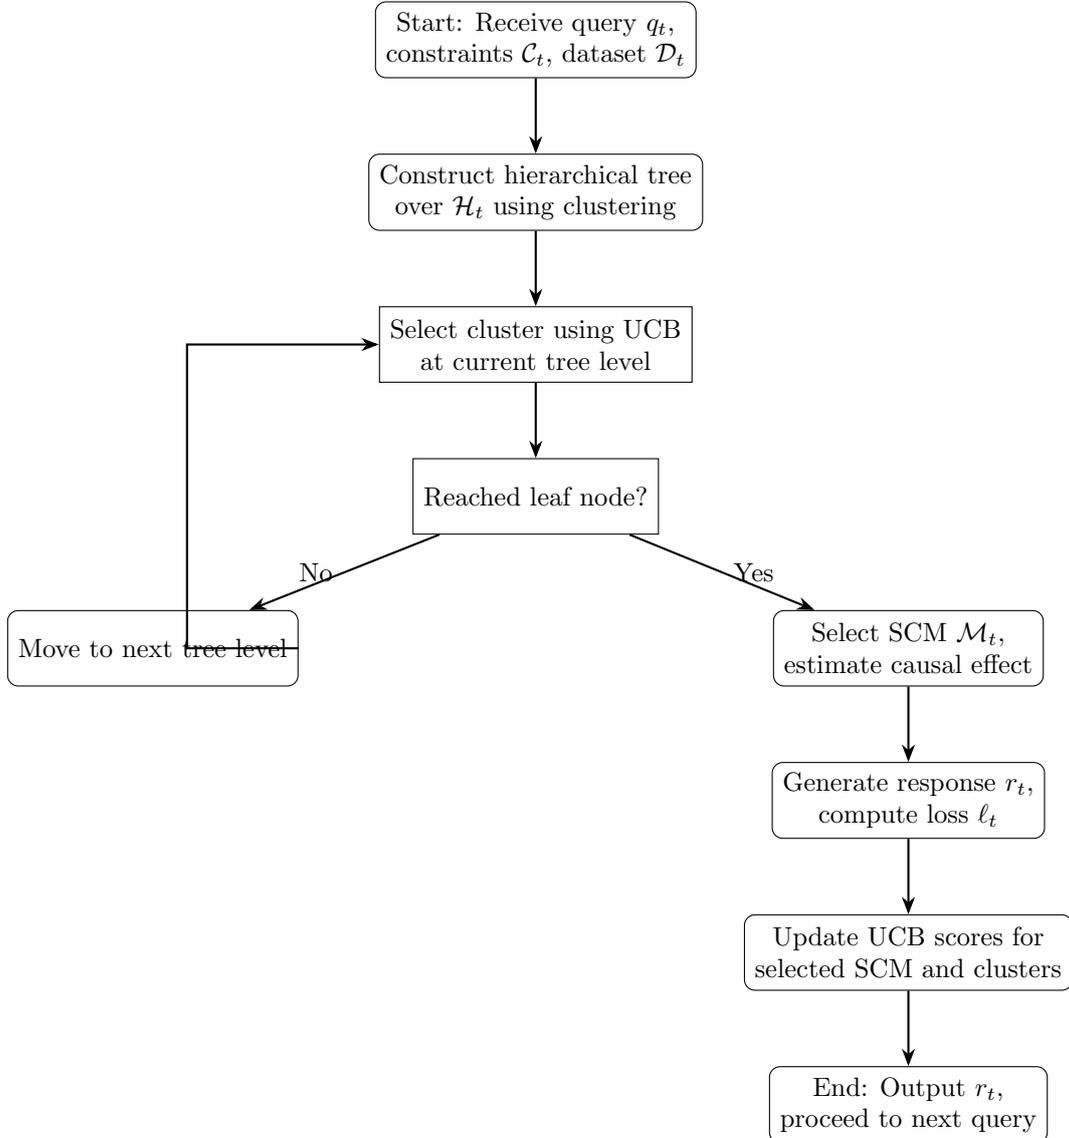


Figure 3: Flowchart of the modified Causal Inferencing algorithm using hierarchical hypothesis spaces, as described in Section A.2. The algorithm constructs a hierarchical tree over the hypothesis space \mathcal{H}_t , selects clusters using the Upper Confidence Bound (UCB) strategy, and estimates causal effects for the chosen SCM \mathcal{M}_t . The process iterates until a response r_t is generated and UCB scores are updated, supporting the regret bound in Proposition 6.