ON MEASURING INFLUENCE IN AVOIDING UNDESIRED FUTURE

Anonymous authors
Paper under double-blind review

ABSTRACT

When a predictive model anticipates an undesired future event, a question arises: what can we do to avoid it? Resolving this forward-looking challenge requires determining the variables that positively influence the future, moving beyond statistical correlations typically exploited for prediction. In this paper, we introduce a novel framework for evaluating the influence of actionable variables in successfully avoiding the undesired future. We quantify influence as the degree to which the probability of success can be increased by altering variables based on the principle of maximum expected utility. While closely related to causal effects, our analysis reveals a counterintuitive insight: influential variables may not necessarily be those with intrinsically strong causal effects on the target. In fact, due to the dynamics of the decision process, it can be highly beneficial to alter a weak causal factor, or even a variable that is not an intrinsic factor at all. We provide a practical implementation for computing the proposed quantity using observational data and demonstrate its utility through empirical studies on synthetic and real-world applications.

1 INTRODUCTION

When an intelligent machine receives a warning from a powerful predictive model anticipating that an undesired event is going to happen, an important question naturally arises: what can be done to avoid this potential future? This is known as the *avoiding undesired future* (AUF) problem (Zhou, 2022), sparking a transition from passively predicting results to proactively influencing them.

Addressing the AUF problem requires determining the variables that can be properly altered to shape a more desirable future. While statistically correlated variables are effectively exploited by modern machine learning (ML) techniques for predicting target variables (Jumper et al., 2021; Achiam et al., 2023; Price et al., 2025), these correlations are often unreliable for influencing the future target. For instance, although ice cream sales and drowning incidents are highly correlated in the summer, suppressing ice cream sales would obviously not prevent drownings, as their superficial correlation arises from a common cause: hot weather. This implies that a general understanding of the underlying mechanisms connecting variables would be essential for settling the AUF problem.

To this end, an intuitive way is to exploit causal variables of the target. Rich tools for discovering causal relations have been developed in the literature (Pearl, 2009; Peters et al., 2017). Nevertheless, the fact that a variable is a cause of the target variable does not imply that altering it will be influential. For example, while a city's reliance on public transportation might be a cause of lengthy commute times, a policy encouraging the use of private cars could fail to save time due to offsetting effects: the positive impact on shortening commute times obtained via private cars could be neutralized by the negative aspect, such as the worsening traffic congestion caused by many more cars on the road. This seems to suggest shifting our attention to variables with non-negligible average causal effects. However, this strategy is also incomplete. As illustrated by the simple case of two actionable variables in Figure 1, it can be highly beneficial to alter a variable with a negligible causal effect. Therefore, a more principled way is needed to properly address the AUF problem.

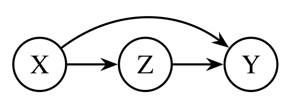

Figure 1: *To do, or not to do, that is the question*: Whether a variable should be altered to influence the eventuality? Let X be a Bernoulli variable, with $Z := 1 - X$ and $Y := X \cdot Z$. Let both X and Z be actionable. Interestingly, while the average causal effect of X on Y is 0, it remains beneficial to alter X. More details are provided in Example 3.

In this paper, we introduce a novel framework for measuring the influence of actionable variables for successfully avoiding the undesired future. We begin by outlining several natural and intuitive considerations that a measure of influence in AUF scenarios should incorporate. Then, we formulate a novel quantity, termed *influence power*, defined as the degree to which the probability of success can be increased through alteration based on the principle of maximum expected utility. This quantity captures the dynamics of decision-making by accounting for the *actionability*, *naturality*, and *desirability* of variables throughout the decision process. In this way, it offers a holistic assessment of the consequence of alteration, capturing both its explicit and implicit impacts on the future target.

Next, we leverage the influence power to investigate the relationship between influential variables and those with an intrinsic causal connection to the future target. Our analysis reveals a subtle yet important nuance: influential variables are not simply a subset of causal ancestors, and vice versa. Specifically, while the influence power is closely related to causal effects, we find that influential variables are not necessarily those with intrinsically strong causal effects. In fact, due to the dynamics of the decision process, it can be highly beneficial to alter a causal ancestor with negligible effects, or even a variable that is not an intrinsic ancestor at all. Another important observation is that, not all actionable variables can be safely altered, as for certain variables, any alteration is counterproductive. This insight crystallizes the fundamental question for an intelligent agent facing an undesired future: *To do, or not to do?* Our framework rests on a principled quantity for measuring influence in AUF, thereby providing a rigorous way to answer this question.

Finally, we address the practical computation of influence power. We identify the challenges inherent in its exact computation and present a Monte-Carlo-based approximation method to efficiently assess it using observational data. The proposed method mitigates the need for full knowledge of structural equations under the assumption of causal sufficiency and tends to remain useful when the probability terms within our quantity are not approximated very perfectly. Empirical studies demonstrate the utility of our framework for addressing the AUF problem on synthetic and real-world applications.

2 PRELIMINARY

Notation. We represent each random variable with a capital letter (V), and its realized value with the lowercase letter (v). We use bold capital letters ($\mathbf{V}$) to denote a set of random variables with their realized values denoted by bold lowercase letters ($\mathbf{v}$). Let $G = (\mathbf{V}, \mathbf{E})$ denote a directed graph with nodes $\mathbf{V}$ and edges $\mathbf{E}$. In a causal graph G, a variable X is a causal ancestor of Y, denoted by $X \in \mathrm{Anc}(Y)$, if there is a directed path from X to Y in G. When X is binary, its causal strength can be quantified by the *average causal effect* (ACE) (Holland, 1988; Pearl, 2009), defined as $\tau(X,Y) := \mathbb{E}(Y|do(X=1) - \mathbb{E}(Y|do(X=0)))$, where $\mathbb{E}(Y|do(X=x))$ denotes the expectation of Y when X is set to the value x. We say that a causal ancestor X of Y is weak if the average causal effect of X on Y is zero. Let Δ_X denote the feasible domain of alteration for a variable X. If $\Delta_X \neq \emptyset$, we call X an actionable variable; otherwise, X is unactionable.

Structural causal models. We use the language of the *structural causal model* (SCM) (Pearl, 2009), which describes how *nature* assigns values to variables of interest, i.e., the physical mechanisms governing the natural generation process of random variables. An SCM is a tuple $\mathcal{M} = \langle \mathbf{V}, \mathbf{N}, F, P(\mathbf{N}) \rangle$, where $\mathbf{V} = \{V_1, \dots, V_d\}$ is a set of endogenous variables, $\mathbf{N} = \{N_1, \dots, N_d\}$ is a set of independent background noises destributed according to $P(\mathbf{N})$, and F is a set of deterministic functions f_i for each $V_i \in \mathbf{V}$ such that $V_i := f_i(\mathrm{PA}_i, N_i)$ with $\mathrm{PA}_i \subseteq \mathbf{V}$. Throughout this paper, we posit that the natural generation process is governed by an underlying SCM $\mathcal{M}$, though it may remain unknown to the decision-maker due to its unobserved nature (Bareinboim et al., 2022). If V_i is an ancestor of V_j in the causal graph induced by the underlying SCM $\mathcal{M}$, we say that V_i is an *intrinsic* ancestor of V_j in $\mathcal{M}$. For a variable V_i, if $\Delta_{V_i} \neq \emptyset$, we use the notation $V_i :- v_i$ to indicate that V_i *can* be altered to $v_i \in \Delta_{V_i}$. This operation replaces the structural function of V_i in $\mathcal{M}$ with the constant assignment $V_i := v_i$, and the distribution of $\mathbf{W}$ given that V_i is set to v_i is denoted as $P(\mathbf{W}|V_i :- v_i)$.

Problem definition. We consider a setting where observational data is drawn from a distribution induced by an underlying SCM $\mathcal{M}$. We suppose that this SCM characterizes the natural generation process of a sequence of variables $(V_1, \dots, V_{d+1})$, where the final variable V_{d+1} represents the target variable Y, whose desired domain is specified as $\mathcal{S}$. The variable sequence is pre-specified and is consistent with the underlying causal structure (i.e., variables are causally ordered with respect to $\mathcal{M}$).[1] For simplicity, we assume that all variables are discrete. The goal of decision-making in the AUF problem is to maximize the possibility of Y falling into $\mathcal{S}$ through feasible alterations on the variables $V_1, \ldots, V_d$. Generally, it is convenient to denote by $\mathbf{x}$ a realization of a subset of variables $\mathbf{X} \subset \{V_1, \ldots, V_d\}$, and denote by $\mathbf{Z}$ the set of actionable variables succeeding $\mathbf{X}$, i.e., $\mathbf{Z} = \{V_i \mid t < i \leq d,\, \Delta_{V_i} \neq \emptyset\}$, where $t = \max\{s \mid V_s \in \mathbf{X}\}$. Hence, when an initial observation $\mathbf{X} = \mathbf{x}$ is given, the AUF problem is addressed by altering variables in $\mathbf{Z}$.[2] Notably, we do not assume that all variables preceding $\mathbf{Z}$ are observed when determining alterations to $\mathbf{Z}$; the initial set $\mathbf{X}$ may even be empty. This makes our framework more practical than those requiring full observability.

3 Influence Power

3.1 Motivation

We motivate the considerations that a measure of influence in AUF should incorporate by describing the strategies and limitations of existing approaches for addressing the AUF problem.

A primary strategy is to find a feasible alteration that directly maximizes the probability of Y falling within the desired domain $\mathcal{S}$ (Qin et al., 2023). This straightforward strategy is expressed as:

$$(Z^*, z^*) = \arg\max_{Z \in \mathbf{Z}, z \in \Delta_Z} P(Y \in \mathcal{S} | \mathbf{X} = \mathbf{x}, Z :- z), \tag{1}$$

where $\mathbf{x}$ is the observation of $\mathbf{X}$, and $\mathbf{Z}$ is the set of actionable variables succeeding $\mathbf{X}$. This approach is intuitive and can indeed achieve a better target in many cases, but it overlooks several important considerations. Specifically, Equation (1) only accounts for the straightforward effect of altering a single variable at a time, presuming a "static" future where subsequent variables unfold naturally. Thus, an immediate consequence is that it ignores how multiple variables might combine their effects. A very simple example illustrates this issue. Imagine two binary variables, Z_1 and Z_2, both of which naturally take the value 0 with near certainty, and let $Y := Z_1 \wedge Z_2$. Clearly, altering either variable alone is ineffective. It's only by altering both variables together that we can achieve $Y = 1$. Consequently, when judging the impact of an alteration in AUF scenarios, not only the feasible domain of the alteration itself but also the *actionability* of other variables should be considered.

Given the insight from the example above, the next logical step would be to propose the joint alteration of all actionable variables as a solution. This joint strategy has been adopted in previous work (Qin et al., 2025; Du et al., 2025) with the following formulation:

$$\mathbf{z}^* = \arg\max_{\mathbf{z} \in \Delta_{\mathbf{Z}}} P(Y \in \mathcal{S} | \mathbf{X} = \mathbf{x}, \mathbf{Z} :- \mathbf{z}), \tag{2}$$

where $\Delta_{\mathbf{Z}}$ denotes the Cartesian product of the feasible domains of alteration for all variables in $\mathbf{Z}$, and $\mathbf{Z} :- \mathbf{z}$ denotes the joint alteration of all variables in $\mathbf{Z}$ to the corresponding values in $\mathbf{z}$. This strategy overlooks an important fact: it's often unnecessary to alter all variables. For instance, while both light and water are crucial factors for crop growth, if sunlight is naturally abundant, adding artificial light will have negligible impact on yield. Therefore, when judging the impact of altering a variable in AUF scenarios, we need to consider its *naturality*, i.e., whether it is in a favorable state naturally. Moreover, as we shall see in what follows, certain variables may not only be unnecessary to alter, but could even be counterproductive no matter how they are altered. Thus, a more principled approach is required to determine which actionable variables should be altered.

3.2 Formulation

In this subsection, we formulate a new quantity that measures whether an actionable variable is worth altering in order to influence the future target. To holistically account for the actionability and naturality of variables, as well as the desirability of the target variable in the decision process, our formulation requires a principled way to envision future possibilities after an alteration. The Bellman equation (Bellman, 1957) provides the conceptual foundation for this purpose, but its standard formulation is not immediately applicable to our context. This is because the classical

[1] This ensures that the sequence of variables satisfies the definition of *statistical time* (Pearl, 2009) and thus inherently accommodates a temporal interpretation.

[2] The variables V_s for $s \leq t$ are immutable, as the past cannot be changed. This situates our work within the scope of Level 2 of the ladder of causation (Pearl & Mackenzie, 2018; Bareinboim et al., 2022).

framework is usually built upon a prespecified separation between state and control variables. In the AUF problem, however, every variable V_i in the sequence $(V_1, \ldots, V_d)$ has a dual role: it could be proactively manipulated through alteration or be passively observed as it unfolds naturally.

Drawing inspiration from the Bellman equation and grounding our proposal in the principle of maximum expected utility (Russell & Norvig, 2020), we recursively define the *maximum expected probability* (MEP) of avoiding the undesired future after an alteration or observation. Specifically, for $0 < k < d$, the MEP after *altering* V_k to v_k is given by

$$\mathcal{P}(Y \in \mathcal{S}|V_k :- v_k, \ldots) \coloneqq \max\Big\{ \max_{v_{k+1}\in\Delta_{V_{k+1}}} \mathcal{P}(Y \in \mathcal{S}|V_{k+1} :- v_{k+1}, V_k :- v_k, \ldots), \quad \mathbb{E}_{v_{k+1}\sim P(V_{k+1}|V_k \,:-\, v_k, \ldots)} \mathcal{P}(Y \in \mathcal{S}|V_{k+1} = v_{k+1}, V_k :- v_k, \ldots)\Big\}, \tag{3}$$

where $\mathcal{P}(Y \in \mathcal{S}|V_{k+1} = v_{k+1}, V_k :- v_k, \ldots)$ is interpreted as the MEP after the observation of $V_{k+1} = v_{k+1}$ and the alteration of $V_k :- v_k$, and "$\ldots$" abbreviates any form of alterations and observations that happened before V_k. For $k = d$, the MEP after altering V_k to v_k, $\mathcal{P}(Y \in \mathcal{S}|V_k :- v_k, \ldots)$, simply equals to the *AUF probability*, $P(Y \in \mathcal{S}|V_k :- v_k, \ldots)$. Similarly, for $0 < k < d$, the MEP after *observing* V_j as v_j is given by

$$\mathcal{P}(Y \in \mathcal{S}|V_j = v_j, \ldots) \coloneqq \max\Big\{ \max_{v_{j+1}\in\Delta_{V_{j+1}}} \mathcal{P}(Y \in \mathcal{S}|V_{j+1} :- v_{j+1}, V_j = v_j, \ldots), \quad \mathbb{E}_{v_{j+1}\sim P(V_{j+1}|V_j \,=\, v_j, \ldots)} \mathcal{P}(Y \in \mathcal{S}|V_{j+1} = v_{j+1}, V_j = v_j, \ldots)\Big\}. \tag{4}$$

For $j = d$, the MEP after observing V_j as v_j, $\mathcal{P}(Y \in \mathcal{S}|V_j = v_j, \ldots)$, simply equals to the AUF probability, $P(Y \in \mathcal{S}|V_j = v_j, \ldots)$.

Based on the above recursive definition of MEP, we formulate a quantity called the *influence power*, indicating the ability of an actionable variable to influence the future target.

Definition 1 (Influence Power). *The influence power of an actionable variable V_i on Y is defined as*

$$\dot{p}(V_i, Y) \coloneqq \max_{v_i\in\Delta_{V_i}} \mathcal{P}(Y \in \mathcal{S}|V_i :- v_i) - \mathbb{E}_{v_i\sim P(V_i)}\mathcal{P}(Y \in \mathcal{S}|V_i = v_i).$$

Remark. The influence power of V_i on Y represents the maximum increase in the MEP that can be achieved by optimally altering V_i, compared to the expected MEP when V_i is observed naturally. Consequently, a positive influence power indicates that the alteration is beneficial, while a zero or negative influence power suggests that it is unnecessary or even harmful. By definition, the influence power is bounded within the range of $[-1, 1]$. As Definition 1 recursively follows the principle of maximum expected utility, the influence power can be interpreted as a variant of the Bellman equation. Notably, this notion can be easily extended to a conditional form: e.g., given the observation $\mathbf{X} = \mathbf{x}$, the conditional influence power of $V_i \in \mathbf{Z}$ on Y is given by $\dot{p}(V_i, Y|\mathbf{X} = \mathbf{x}) \coloneqq \max_{v_i\in\Delta_{V_i}} \mathcal{P}(Y \in \mathcal{S}|V_i :- v_i, \mathbf{X} = \mathbf{x}) - \mathbb{E}_{v_i\sim P(V_i|\mathbf{X} \,=\, \mathbf{x})}\mathcal{P}(Y \in \mathcal{S}|V_i = v_i, \mathbf{X} = \mathbf{x})$.

We end this subsection by highlighting a connection between Definition 1 and Equation (1). Consider a scenario with three binary variables: V_1, V_2, and Y, where both V_1 and V_2 are actionable. Suppose an oracle informs us that the structural function f defining the target variable Y depends solely on V_1 and not V_2, i.e., $Y \coloneqq f(V_1)$. Based on this information, we deduce that the solution to Equation (1) is V_1 if the following condition holds:

$$\max_{v_1\in\Delta_{V_1}} P(Y \in \mathcal{S}|V_1 :- v_1) > P(Y \in \mathcal{S}). \tag{5}$$

On the other hand, the influence power of V_1 on Y simplifies to

$$\dot{p}(V_1, Y) = \max_{v_1\in\Delta_{V_1}} P(Y \in \mathcal{S}|V_1 :- v_1) - P(Y \in \mathcal{S}). \tag{6}$$

Combining Equations (5) and (6) concludes that the solution of Equation (1) is V_1 if $\dot{p}(V_1, Y) > 0$. Thus, Equation (1) aligns with Definition 1 in determining whether V_1 should be altered.

Furthermore, let $\Delta_{V_1} = \{0, 1\}$ and $\mathcal{S} = \{1\}$. By applying the identity $2 \cdot \max(a, b) = a + b + |a - b|$, the condition of $\dot{p}(V_1, Y) > 0$ reduces to

$$|\tau(V_1, Y)| \equiv |\mathbb{E}(Y|V_1 :- 1) - \mathbb{E}(Y|V_1 :- 0)| > 2\mathbb{E}(Y) - \mathbb{E}(Y|V_1 :- 0) - \mathbb{E}(Y|V_1 :- 1), \tag{7}$$

where $|\tau(V_1, Y)|$ is the absolute value of average causal effect of V_1 on Y. This reveals that $\dot{p}(V_1, Y)$ is closely related to $\tau(V_1, Y)$, and the influence power seems to favor altering variables with strong causal effects. This view, however, is incomplete. In the following section, we will demonstrate that the relationship between influence power and average causal effect is, in fact, far more nuanced.

3.3 Connection

In this subsection, we investigate the connection between influential variables and those with an intrinsic causal relationship to the target variable. Concretely, we analyze how variables with a non-zero influence power, $\dot{p}(X, Y) \neq 0$, generally relate to the intrinsic ancestors of the target in the underlying SCM, $\mathrm{Anc}(Y)$. We also examine the qualitative relationship between the influence power, $\dot{p}(X, Y)$, and the average causal effect, $\tau(X, Y)$, a widely used measure of causal strength. The results are formally summarized in the following theorem.

Theorem 1. *Let X and Y be two endogenous variables in an SCM. The following statements hold:*

1. $X \in \mathrm{Anc}(Y) \not\Rightarrow \dot{p}(X, Y) \neq 0$ *and* $\dot{p}(X, Y) \neq 0 \not\Rightarrow X \in \mathrm{Anc}(Y)$;
2. $\tau(X, Y) \neq 0 \not\Rightarrow \dot{p}(X, Y) \neq 0$ *and* $\dot{p}(X, Y) \neq 0 \not\Rightarrow \tau(X, Y) \neq 0$;
3. $\tau(X, Y) \neq 0 \not\Rightarrow \dot{p}(X, Y) \geq 0$ *and* $X \in \mathrm{Anc}(Y) \not\Rightarrow \dot{p}(X, Y) \geq 0$.

Theorem 1 reveals an intricate relationship between intrinsic ancestors and non-zero influence power: neither implies the other. Specifically, a causal ancestor of the target can have zero influence power, and conversely, a variable with non-zero influence power is not necessarily an intrinsic ancestor in the underlying SCM. Similarly, a variable may have a non-zero average causal effect on the target while manifesting zero influence power, and vice versa. Furthermore, neither having a non-zero average causal effect nor being a causal ancestor guarantees non-negative influence power. These relationships are visualized as a Venn diagram in Figure 2. In the following, we shed light on several insights of statements in Theorem 1 with concrete examples, while a detailed proof is deferred to Appendix B. To facilitate understanding, these examples focus on binary variables with the desired domain $\mathcal{S} = \{1\}$, though this restriction is not required generally.

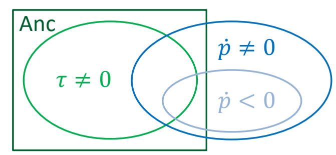

Figure 2: Relationship between intrinsic ancestors, τ, and $\dot{p}$.

A causal ancestor can have zero influence power. Altering a causal ancestor with a strong average causal effect on the target may provide no benefit to the target.

Example 1. *Consider the following structural equations with the corresponding causal graph:*

$$X := N_X,$$
$$Z := X \cdot N_Z + (1 - X) \cdot (1 - N_Z),$$
$$Y := Z \cdot N_Y + (1 - Z) \cdot (1 - N_Y),$$

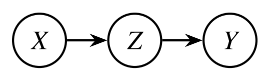

where $N_X, N_Z, N_Y \overset{iid}{\sim} \mathrm{Bern}(0.9)$. *Let X and Z be actionable variables, let $\Delta_X = \{0, 1\}$ and $\Delta_Z = \{0, 1\}$ be the feasible domains of alteration, and let the desired domain for Y be $\mathcal{S} = \{1\}$.*

In this example, while X is an ancestor of Y in the SCM, its influence power on Y is zero: $\dot{p}(X, Y) = \max_{x \in \Delta_X} \mathcal{P}(Y = 1 | X :- x) - \mathbb{E}_{x \sim P(X)} \mathcal{P}(Y = 1 | X = x) = \max_{x \in \Delta_X} \max_{z \in \Delta_Z} P(Y = 1 | Z :- z, X :- x) - \mathbb{E}_{x \sim P(X)} \max_{z \in \Delta_Z} P(Y = 1 | Z :- z, X = x) = 0.9 - 0.9 = 0$. This indicates that altering X yields no improvement in the probability of $Y = 1$; a rational machine will always maximize the probability of $Y = 1$ by setting Z to 1, regardless of the value of X. In a word, altering X in Example 1 is useless as X is shielded by the actionability of Z. Thus, a causal ancestor does not necessarily have non-zero influence power. In addition, the average causal effect of X on Y in the SCM is non-zero: $\tau(X, Y) = P(Y = 1 | X :- 1) - P(Y = 1 | X :- 0) = 0.82 - 0.18 = 0.64$. This shows that a non-zero average causal effect does not guarantee non-zero influence power.

A non-ancestral variable can have non-zero influence power. Altering a variable that is not an intrinsic ancestor of the target in the underlying SCM may still benefit the target.

Example 2. *Consider the following structural equations with the corresponding causal graph:*

$$U := N_U,$$
$$W := N_W,$$
$$X := U \cdot W \cdot (1 - N_X),$$
$$Z := N_Z,$$
$$Y := Z \cdot (1 - U) + (1 - Z) \cdot N_Y,$$

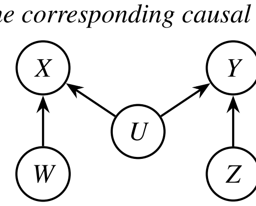

where $N_U \sim \text{Bern}(0.5)$*,* $N_W, N_X, N_Z \overset{iid}{\sim} \text{Bern}(0.1)$*, and* $N_Y \sim \text{Bern}(0.4)$*. Let* W*,* X*, and* Z *be actionable variables with* $\Delta_W = \Delta_X = \Delta_Z = \{0, 1\}$*, and let the desired domain be* $\mathcal{S} = \{1\}$*.*

In this example, W is not an ancestor of Y in the SCM. Remarkably, the influence power of W on Y is positive: $\dot{p}(W, Y) = 0.68 - 0.518 = 0.162$. This indicates that altering W can significantly improve the MEP of $Y = 1$. Intuitively, this positive influence manifests because altering W could help X to reveal information about U, which facilitates a more informed alteration on Z, ultimately benefiting Y. To provide concrete intuition, let us ground the variables from Example 4 in a medical scenario: let U, W, X, Z, and Y denote an allergy gene, a skin test, the skin response, a drug injection, and patient recovery, respectively. Performing a skin test (W) has no therapeutic effect; thus, the average causal effect of W on Y is zero. Nevertheless, the skin test is crucial because it informs the doctor's decision on administering the drug (Z), which impacts recovery (Y). For instance, if the skin test is positive (observing $X = 1$ after setting $W :- 1$), the doctor the doctor can infer the presence of the allergy and decide not to administer the drug ($Z :- 0$), thereby maximizing the probability of recovery ($Y = 1$). This shows that while W does not intrinsically cause Y, altering W is instrumental for positively influencing Y. Influence power successfully captures this implicit benefit, showing that even non-ancestral variables can be critical for AUF.

For completeness, we also examine the conditional influence power of W on Y given U. We find that $\dot{p}(W, Y|U = 1) = 0$ and $\dot{p}(W, Y|U = 0) = 0$. This implies that if the allergy gene (U) were observed, performing the skin test (W) would be unnecessary. In clinical practice, however, directly observing the allergy gene (U) for a new patient is often time-consuming or prohibitively expensive. Thus, the unconditional influence power remains instructive for addressing the AUF problem.

A weak ancestor can have positive influence power. Altering a causal ancestor with a negligible average causal effect on the target may still benefit the target.

Example 3. *Consider the following structural equations with the corresponding causal graph:*

$$\begin{aligned} X &\coloneqq N_X, \\ Z &\coloneqq (1 - X) \cdot N_Z, \\ Y &\coloneqq X \cdot Z \cdot N_Y, \end{aligned}$$

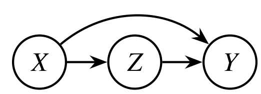

where $N_X, N_Z, N_Y \overset{iid}{\sim} \text{Bern}(0.5)$*. Let* X *and* Z *be actionable variables with* $\Delta_X = \{0, 1\}$ *and* $\Delta_Z = \{0, 1\}$*, and let the desired domain for* Y *be* $\mathcal{S} = \{1\}$*.*

In this example, X is an ancestor of Y, and the average causal effect is zero: $\tau(X, Y) = 0$. Yet, the influence power of X on Y is positive: $\dot{p}(X, Y) = 0.25$. Intuitively, this positive influence power manifests from the synergy between X and Z. The benefit of X on Y is elicited when we account for the alteration of Z. This implicit impact is captured by $\dot{p}(X, Y)$ but missed by $\tau(X, Y)$.

A strong ancestor can have negative influence power. Altering a causal ancestor with a non-negligible average causal effect can be not only useless but also detrimental to the target.

Example 4. *Consider the following structural equations with the corresponding causal graph:*

$$\begin{aligned} U &\coloneqq N_U, \\ X &\coloneqq U \cdot N_X + (1 - U) \cdot (1 - N_X), \\ Z &\coloneqq X \cdot N_Z + (1 - X) \cdot (1 - N_Z), \\ Y &\coloneqq Z \cdot (1 - U) + (1 - Z) \cdot N_Y, \end{aligned}$$

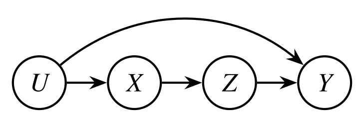

where $N_U, \sim \text{Bern}(0.5)$*,* $N_X, N_Z \overset{iid}{\sim} \text{Bern}(0.9)$*, and* $N_Y \sim \text{Bern}(0.4)$*. Let* X *and* Z *be actionable variables with* $\Delta_X = \{0, 1\}$ *and* $\Delta_Z = \{0, 1\}$*, and let the desired domain be* $\mathcal{S} = \{1\}$*.*

In this example, X is an ancestor of Y in the SCM with a non-zero average causal effect: $\tau(X, Y) = 0.08$, whereas the influence power is negative: $\dot{p}(X, Y) = -0.15$. This indicates that the MEP after altering X is lower than the expected MEP after observing X. Thus, any alteration on X is counterproductive regardless of the specific value to which X is set. Intuitively, this negative influence manifests because observing X reveals information about U, which is useful in determining the alteration on Z during the computation of $\dot{p}(X, Y)$. Hence, while altering X can produce a straightforward improvement in Y (as indicated by the non-zero $\tau(X, Y)$), this benefit is overturned by the negative consequence for the alteration of the subsequent variable, ultimately making the alteration of X detrimental. Again, this implicit impact is successfully captured by $\dot{p}(X, Y)$.

4 ESTIMATING INFLUENCE POWER

Influence power is a principled quantity for measuring the influence of actionable variables in AUF, while computing it exactly is often intractable due to the need for an exhaustive computation of the MEP terms and full knowledge of the underlying structural equations. In this section, we present a practical estimation method to mitigate these challenges.

4.1 MONTE-CARLO APPROXIMATION

The recursive enumeration of MEP for all possible alterations can be computationally prohibitive when the number of actionable variables is large. To mitigate this, we interpret the computation of MEP as a single-player non-deterministic game and approximate it based on the Monte-Carlo tree search UCT (Upper Confidence Tree) introduced by Kocsis & Szepesvári (2006).

Specifically, a *search tree* employing Monte-Carlo simulations is constructed incrementally. Each *node* in the tree represents a *state* defined by a sequence of alterations and observations made so far, associated with the next variable to be considered. Every iteration begins at the root node N_0 (associated with a pre-specified variable $V_i \in \mathbf{V}$), proceeds to its children (associated with V_{i+1}), and continues until reaching a terminal state (associated with the target variable Y). Each *edge* in the tree represents a *choice* that can be made from the node, i.e., either an alteration or an observation on the associated variable. The overall construction consists of four steps, iterated until time has expired: (1) *Selection*: starting from the root node, recursively select an edge to child nodes according to the UCT policy until reaching a leaf node; (2) *Expansion*: if the leaf node corresponds to a non-terminal state, expand it by randomly adding one child node corresponding to possible choices; (3) *Playout*: from the newly added node, execute a random sequence of choices until reaching a terminal state, and compute the AUF probability at that terminal state; (4) *Backpropagation*: propagate the computed AUF probability back up the tree, updating the statistics of each node along the path. During each iteration, the UCT criterion is used at a node N to select the next edge to traverse:

$$c_N^* = \arg\max_{c \in \Delta_N^+} \left\{ \hat{p}_{N,c} + \alpha \cdot \sqrt{\frac{\ln t_N}{t_{N,c}}} \right\}, \tag{8}$$

where $\Delta_N^+ = \Delta_N \cup \emptyset$ is the set of choices at node N (comprising feasible alterations on the variable associated with N, denoted by Δ_N, and the option to make an observation, denoted by $\emptyset$), $\hat{p}_{N,c}$ is the average AUF probability obtained after taking choice c at node N, α is a parameter used to balance between exploration and exploitation (Auer et al., 2002), t_N is the number of times node N has been selected, and $t_{N,c}$ is the number of times choice c has been selected at node N.

After the construction of search tree, the MEP terms in the influence power of V_i on Y are approximated as the average AUF probability for each choice at the root node N_0 of search tree. Concretely, we have $\mathcal{P}(Y \in \mathcal{S} | V_i :- c) \approx \hat{p}_{N_0,c}$ for each $c \in \Delta_{N_0}$, and $\mathbb{E}_{v_i \sim P(V_i)} \mathcal{P}(Y \in \mathcal{S} | V_i = v_i) \approx \hat{p}_{N_0,\emptyset}$. Hence, according to Definition 1, the influence power of V_i on Y is approximated as

$$\dot{p}(V_i, Y) \approx \max_{c \in \Delta_{N_0}} \hat{p}_{N_0,c} - \hat{p}_{N_0,\emptyset}. \tag{9}$$

The quality of this approximation improves over time, as UCT is guaranteed to converge to the best choice given sufficient iterations. Moreover, the described procedure is an *anytime* algorithm, capable of producing an approximate influence power at any point during its computation. We refer the reader to Browne et al. (2012) for further details.

Finally, we note that Equation (9) can remain a useful indicator with a limited number of Monte-Carlo simulations. This is because that a highly accurate estimate of influence power is not always necessary for the AUF problem; in many cases, a rough approximation is enough. Specifically, if the ground-truth influence power of a variable is non-positive ($\dot{p} \leq 0$), the approximation succeeds as long as it correctly suggests that no alteration on the variable is beneficial. This simply requires the approximated MEP terms to satisfy $\hat{p}_{N_0,\emptyset} \geq \max_{c \in \Delta_{N_0}} \hat{p}_{N_0,c}$. Similarly, if the ground-truth influence power is positive ($\dot{p} > 0$), the approximation succeeds as long as it accurately identifies the optimal alteration c^*, which implies that the relative magnitude of the MEP terms is correct. This only requires the approximated MEP terms to satisfy $\hat{p}_{N_0,c^*} \geq \max_{c \in \Delta_{N_0}} \hat{p}_{N_0,c}$. Thus, even with imperfect approximation, the method can still provide reliable indications of influence for AUF.

4.2 AUF Probability Estimation

Although the Monte-Carlo procedure described above can effectively approximate the influence power, it still relies on the AUF probability when a terminal state is reached during simulations, whose ground-truth value is dictated by the underlying SCM. For situations where the structural equations are unknown, we present an expression for estimating the AUF probability from observational data.

Specifically, we express the joint probability of the ordered variables $(\mathbf{V}, Y)$ as:

$$P(\mathbf{V}, Y) = P(V_1, \ldots, V_d, Y) = P(Y|\mathbf{V}) \prod_{i=1}^{d} P(V_i|V_1, \ldots, V_{i-1}), \tag{10}$$

where the conditional probabilities $P(Y|\mathbf{V})$ and $P(V_i|V_1, \ldots, V_{i-1})$ can be estimated from observational data $\mathcal{D} = \{(\mathbf{v}^j, y^j)\}_{j=1}^n$ using standard ML techniques. Denote by $\mathbf{A}$ the variables in $\mathbf{V}$ that are altered, we express the joint probability of $(\mathbf{V}, Y)$ given the alteration of $\mathbf{A}$ as follows:

$$P(\mathbf{V}, Y|\hat{\mathbf{A}}) = P(Y|\mathbf{V}) \prod_{V_i \in \mathbf{A}} \delta(V_i) \prod_{V_i \in \mathbf{V} \setminus \mathbf{A}} P(V_i|V_1, \ldots, V_{i-1}), \tag{11}$$

where $\hat{\mathbf{A}}$ indicates that every variable $V_i \in \mathbf{A}$ is altered, and $\delta(\cdot)$ is the Dirac delta function. Then, denote by $\mathbf{O}$ the variables in $\mathbf{V}$ that are observed, the AUF probability given the alteration of $\mathbf{A}$ and the observation of $\mathbf{O}$ is expressed as:

$$\begin{aligned} P(Y \in \mathcal{S}|\hat{\mathbf{A}}, \mathbf{O}) &= \frac{P(Y \in \mathcal{S}, \mathbf{O}|\hat{\mathbf{A}})}{P(\mathbf{O}|\hat{\mathbf{A}})} = \frac{\sum_{\mathbf{V} \setminus \mathbf{O}} P(Y \in \mathcal{S}, \mathbf{V}|\hat{\mathbf{A}})}{\sum_{\mathbf{V} \setminus \mathbf{O}} P(\mathbf{V}|\hat{\mathbf{A}})} \\ &= \frac{\sum_{\mathbf{V} \setminus \mathbf{O}} P(Y \in \mathcal{S}|\mathbf{V}) \prod_{V_i \in \mathbf{A}} \delta(V_i) \prod_{V_i \in \mathbf{V} \setminus \mathbf{A}} P(V_i|V_1, \ldots, V_{i-1})}{\sum_{\mathbf{V} \setminus \mathbf{O}} \prod_{V_i \in \mathbf{A}} \delta(V_i) \prod_{V_i \in \mathbf{V} \setminus \mathbf{A}} P(V_i|V_1, \ldots, V_{i-1})}, \end{aligned} \tag{12}$$

which is a generic expression of the AUF probability given any alterations and observations. It can be estimated from observational data $\mathcal{D}$ and then plugged into the Monte-Carlo procedure described above to approximate the influence power. The following proposition demonstrates the consistency of Equation (12) by leveraging the manipulation theorem in Spirtes et al. (2000).

Proposition 1. *Assume causal sufficiency, i.e., the joint distribution $P(\mathbf{V}, Y)$ is induced by an acyclic SCM $\mathcal{M}$ with mutually independent background noises, and positivity, i.e., $P(V_i|\mathrm{PA}_i) > 0$ in the support of P, $\forall 1 \leq i \leq d$. Then, the expression in Equation* (11) *is consistent to the joint probability dictated by the SCM $\mathcal{M}_{\mathbf{A}}$ where variables $\mathbf{A}$ are altered. Furthermore, the expression in Equation* (12) *is consistent to the AUF probability dictated by the SCM $\mathcal{M}_{\mathbf{A}}$ where variables $\mathbf{A}$ are altered and variables $\mathbf{O}$ are observed.*

Remark. Causal sufficiency is required in Proposition 1 but is not assumed for the rest of the paper. Technically, this assumption is important for reliably estimating the conditional probabilities in Equations (11) and (12) from observational data. Once these probabilities are estimated, some variables can remain unobserved when deciding whether to alter a variable. This decoupling is more practical than assuming full observability throughout both the estimation and decision phases.

5 Experiments

In this section, we conduct experiments to validate the utility of our framework.

Tasks. We simulate three synthetic tasks (including Trader, Farmer, and Doctor) and a real-world case study (Bermuda). For each task, we generate 1000 samples from the underlying SCM to form the observational data and repeat the experiments ten times. The details of the tasks are provided in Appendix A due to space limitation.

Baselines. We compare six methods for selecting alterations: (1) Observe: a baseline that only observes without altering variables; (2) Max-One: selects the single variable with the highest AUF probability for alteration, as described in Equation (1); (3) Max-All: selects all actionable variables for alteration, as described in Equation (2); (4) MIS: alters a variable if it belongs to the minimal intervention set defined in Lee & Bareinboim (2018); (5) VoC: alters a variable when doing so increases the AUF probability of altering the next variable (Everitt et al., 2021); and (6) Ours: uses MCTS to perform 1000 iterations to determine whether and how to alter variables based on influence power. The parameter α is set to $\sqrt{2}$ by default following Kocsis & Szepesvári (2006). For fair

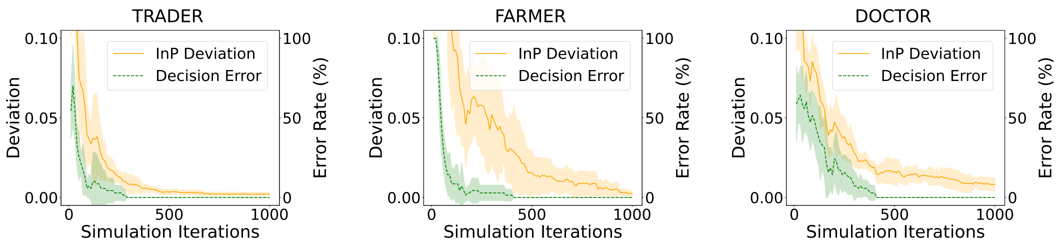

Figure 3: Convergence of the approximation of influence power (InP) and error rates (%) versus the number of MCTS iterations. The deviation of the approximated influence power to the exact value of influence power continues to decrease after the convergence of error rates in all cases.

Task	Observe	Max-One	Max-All	MIS	VoC	Ours
Trader	38.34 ± 3.69	50.22 ± 5.90	50.50 ± 5.69	50.55 ± 7.36	53.20 ± 7.27	62.11 ± 9.05
Farmer	11.03 ± 4.66	56.32 ± 13.89	56.60 ± 12.98	56.70 ± 13.59	57.17 ± 13.86	57.94 ± 12.54
Doctor	39.47 ± 4.87	50.83 ± 5.36	50.64 ± 5.32	51.31 ± 6.41	53.72 ± 4.05	65.69 ± 8.06

Table 1: Success rates (%) of six different methods for three synthetic tasks.

Task	10	50	100	500	1000	5000
Trader	43.42 ± 5.97	50.06 ± 6.11	49.27 ± 9.59	59.48 ± 10.70	62.11 ± 9.05	62.39 ± 9.37
Farmer	20.21 ± 10.29	31.39 ± 9.83	53.32 ± 14.24	56.54 ± 14.05	57.94 ± 12.54	56.46 ± 14.00
Doctor	43.86 ± 4.78	42.56 ± 3.99	48.58 ± 7.29	64.48 ± 9.60	65.69 ± 8.06	66.26 ± 6.66

Table 2: Success rates (%) of our method with different sample sizes for three synthetic tasks.

Task	0.01	0.1	0.5	1.0	2.0	10.0
Trader	55.84 ± 7.69	58.08 ± 7.45	60.15 ± 4.37	61.04 ± 5.48	62.73 ± 5.40	59.04 ± 3.92
Farmer	55.79 ± 11.46	55.95 ± 14.66	57.93 ± 13.46	57.36 ± 12.75	56.22 ± 13.36	55.50 ± 13.86
Doctor	58.31 ± 5.93	59.90 ± 5.68	66.75 ± 7.14	65.92 ± 9.06	65.60 ± 6.11	58.22 ± 5.36

Table 3: Success rates (%) of our method with different α for three synthetic tasks.

comparison, the feasible domain for each actionable variable is set to be $\{0, 1\}$ and the number of actionable variables is set to 3 for all methods in synthetic tasks. The performance of each method is evaluated by the success rate, i.e., the frequency of the target variable successfully achieving the desired domain after performing alterations on the suggested variables.

Figure 3 shows the convergence of approximating influence power. The plot depicts the deviation of the approximated value for the first actionable variable, measured as the absolute difference from the corresponding exact value. The error rate represents the frequency of inconsistencies between the suggested alterations based on the approximated value and the exact value. In all cases, the error rate decrease as T increases, demonstrating the effectiveness of MCTS in approximating influence power. Notably, the deviation continues to decrease after the error rate has converged to zero, demosntrating that our method is useful when the MEP terms are not approximated very perfectly.

Table 1 compares our method with baselines. We observe that our method consistently outperforms existing methods in most cases. These results demonstrate the superiority of the proposed method in guiding alterations for AUF tasks. In the Farmer task, various methods perform comparably. This is because the target variable in this specific task is influenced by a single critical variable, which all five methods correctly determined. Table 2 investigates the impact of sample size on the effectiveness of our method. The performance generally improves as the sample size increases. Notably, the success rates exhibit a rapid growth initially and begin to plateau, stabilizing around 1,000 samples across the tasks. The sensitivity of the hyperparameter α is reported in Table 3. The results indicate robust performance for values between 0.5 and 2.0, where the method achieves consistently high success rates. Extreme values (too small, e.g., 0.01, or too large, e.g., 10.0) degrade performance.

TASK	OBSERVE	MAX-ONE	MAX-ALL	MIS	VOC	OURS
BERMUDA	2.36 ± 0.50	61.22 ± 0.91	72.68 ± 2.60	75.06 ± 1.67	63.44 ± 0.37	78.45 ± 0.56

Table 4: Success rates (%) of six different methods for the BERMUDA task.

Furthermore, we experiment with BERMUDA, a real-world application to further evaluate the utility of our method. As presented in Table 4, our method achieves the highest performance with a success rate of 78.45%, surpassing the second-best baseline by a clear margin. These results demonstrate that our approach remains robust and effective in a complex real-world scenario involving non-binary variables, where it consistently outperforms existing methods.

6 RELATED WORK

Many efforts have been dedicated to identify causal structures and causal effects from observational data in the literature (Verma & Pearl, 1991; Cooper & Herskovits, 1992; Heckerman et al., 1995; Zheng et al., 2018; Lorch et al., 2021). Apart from the average causal effect (Rosenbaum & Rubin, 1983; Holland, 1988), there are various other quantities for measuring causal strength such as analysis of variance (Northcott, 2008) or other approaches (Janzing et al., 2013; Jung et al., 2022). We primarily focus on comparing with average causal effects, as it is a popular and canonical measure of causal strength in the literature. The comparison regarding other measures of causal strength would be similar and left for future work. We also note that researchers have proposed various ways of quantifying the strength of causal contributions, sometimes referred to as "causal influence" (Rosenbaum & Rubin, 1983; Holland, 1988; Janzing et al., 2013; Heskes et al., 2020). Different notions of influence coexist for good reason, as they formalize different perspectives on different goals (Janzing et al., 2024). Much of the prior work has focused on quantifying intrinsic causal contributions, i.e., the degree to which various factors "explain" the variance of a target variable, which is valuable for attribution and scientific understanding. This work, in contrast, focuses on quantifying practical utility for decision-making in the AUF problem.

This work is essentially distinct from approaches based on counterfactual reasoning (Pearl, 2009; Halpern, 2015; Karimi et al., 2021; Tsirtsis et al., 2021). While counterfactuals generally involve reasoning about the past (i.e., what would have happened, had we chosen differently at a point in the past (Pearl et al., 2016)), the AUF problem is forward-looking (i.e., planning for the future). Although some approaches share connections with the Bellman equation (Zhang & Bareinboim, 2019; Tsirtsis et al., 2021), they differ in objective and formalization. Specifically, unlike methods that maintain a strict distinction between state and action variables, our framework treats all variables uniformly as random variables; one can choose to alter (set the value) or explicitly refrain from altering (letting it occur naturally). In addition, we estimate AUF probabilities from observational data without assuming a known causal structure. Besides, compared to concepts like value of control (VoC) (Everitt et al., 2021), which typically assumes that a decision node has no "natural value" and restricts decisions to a single node, our formulation accommodates an arbitrary number of actionable variables with natural generation processes. Additional related works are discussed in Appendix D.

7 CONCLUSION

In this paper, we aim to measure the influence of actionable variables in avoiding the undesired future. Drawing on intuitive considerations, we introduce a novel quantity called influence power, designed to evaluate the extent to which variables can be manipulated in increasing the AUF probability under the principle of maximum expected utility. While closely related to causal effects, our analysis reveals a counterintuitive insight that non-ancestral variables can have non-trivial influence power on the future target. We further provide a practical implementation based on a Monte Carlo-based method to estimate the probability terms in the proposed quantity using observational data, facilitating the efficient approximation of influence power. Experiments on synthetic and real-world tasks validate the utility of our framework in suggesting alterations for addressing the AUF problem.

REPRODUCIBILITY STATEMENT

To ensure the reproducibility of our work, we provide detailed references to experimental setups and theoretical assumptions. Regarding the empirical evaluation, the experimental setups are outlined in Section 5, while comprehensive details, including task specifications, data generation processes, and exact hyperparameter configurations, are documented in Appendix A. The implementation leverages the DOWHY library (Sharma & Kiciman, 2020; Blöbaum et al., 2024), and the code to reproduce our results will be made publicly available upon publication. On the theoretical side, the problem definitions and assumptions are clarified in Section 2 and the remark following Proposition 1, with complete proofs for all claims provided in Appendix C.

REFERENCES

Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*, 2023.

Virginia Aglietti, Xiaoyu Lu, Andrei Paleyes, and Javier González. Causal bayesian optimization. In *International Conference on Artificial Intelligence and Statistics*, pp. 3155–3164, 2020.

Andreas Andersson and Nicholas Bates. In situ measurements used for coral and reef-scale calcification structural equation modeling including environmental and chemical measurements, and coral calcification rates in bermuda from 2010 to 2012 (beacon project), 2018, 2018.

Peter Auer, Nicolo Cesa-Bianchi, and Paul Fischer. Finite-time analysis of the multiarmed bandit problem. *Machine learning*, 47(2):235–256, 2002.

Elias Bareinboim, Juan D Correa, Duligur Ibeling, and Thomas Icard. On pearl's hierarchy and the foundations of causal inference. In *Probabilistic and causal inference: the works of judea pearl*, pp. 507–556. 2022.

Richard Bellman. *Dynamic Programming*. Princeton University Press, Princeton, 1957.

Patrick Blöbaum, Peter Götz, Kailash Budhathoki, Atalanti A. Mastakouri, and Dominik Janzing. Dowhy-gcm: An extension of dowhy for causal inference in graphical causal models. *Journal of Machine Learning Research*, 25(147):1–7, 2024. URL http://jmlr.org/papers/v25/22-1258.html.

Cameron B Browne, Edward Powley, Daniel Whitehouse, Simon M Lucas, Peter I Cowling, Philipp Rohlfshagen, Stephen Tavener, Diego Perez, Spyridon Samothrakis, and Simon Colton. A survey of monte carlo tree search methods. *IEEE Transactions on Computational Intelligence and AI in games*, 4(1):1–43, 2012.

Gregory F Cooper and Edward Herskovits. A bayesian method for the induction of probabilistic networks from data. *Machine learning*, 9:309–347, 1992.

Travis A Courtney, Mario Lebrato, Nicholas R Bates, Andrew Collins, Samantha J De Putron, Rebecca Garley, Rod Johnson, Juan-Carlos Molinero, Timothy J Noyes, Christopher L Sabine, et al. Environmental controls on modern scleractinian coral and reef-scale calcification. *Science advances*, 3(11):e1701356, 2017.

James E Driskell, Carolyn Copper, and Aidan Moran. Does mental practice enhance performance? *Journal of applied psychology*, 79(4):481, 1994.

Wen-Bo Du, Tian Qin, Tian-Zuo Wang, and Zhi-Hua Zhou. Avoiding undesired future with minimal cost in non-stationary environments. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024.

Wen-Bo Du, Hao-Yi Lei Lei, Lue Tao, Tian-Zuo Wang, and Zhi-Hua Zhou. Enabling optimal decisions in rehearsal learning under care condition. In *International Conference on Machine Learning*, 2025.

Tom Everitt, Ryan Carey, Eric D Langlois, Pedro A Ortega, and Shane Legg. Agent incentives: A causal perspective. In *Proceedings of the AAAI conference on artificial intelligence*, volume 35, pp. 11487–11495, 2021.

Joseph Y Halpern. A modification of the halpern-pearl definition of causality. In *Proceedings of the 24th International Conference on Artificial Intelligence*, pp. 3022–3033, 2015.

David Heckerman, Dan Geiger, and David M Chickering. Learning bayesian networks: The combination of knowledge and statistical data. *Machine learning*, 20:197–243, 1995.

Tom Heskes, Evi Sijben, Ioan Gabriel Bucur, and Tom Claassen. Causal shapley values: Exploiting causal knowledge to explain individual predictions of complex models. In *Advances in neural information processing systems*, volume 33, pp. 4778–4789, 2020.

Paul W Holland. Causal inference, path analysis and recursive structural equations models. *ETS Research Report Series*, 1988(1):i–50, 1988.

Dominik Janzing, David Balduzzi, Moritz Grosse-Wentrup, and Bernhard Schölkopf. Quantifying causal influences. *The Annals of Statistics*, 41(5):2324–2358, 2013.

Dominik Janzing, Patrick Blöbaum, Atalanti A Mastakouri, Philipp M Faller, Lenon Minorics, and Kailash Budhathoki. Quantifying intrinsic causal contributions via structure preserving interventions. In *International Conference on Artificial Intelligence and Statistics*, pp. 2188–2196. PMLR, 2024.

John Jumper, Richard Evans, Alexander Pritzel, Tim Green, Michael Figurnov, Olaf Ronneberger, Kathryn Tunyasuvunakool, Russ Bates, Augustin Žídek, Anna Potapenko, et al. Highly accurate protein structure prediction with alphafold. *nature*, 596(7873):583–589, 2021.

Yonghan Jung, Shiva Kasiviswanathan, Jin Tian, Dominik Janzing, Patrick Blöbaum, and Elias Bareinboim. On measuring causal contributions via do-interventions. In *International Conference on Machine Learning*, pp. 10476–10501. PMLR, 2022.

Amir-Hossein Karimi, Bernhard Schölkopf, and Isabel Valera. Algorithmic recourse: from counterfactual explanations to interventions. In *Proceedings of the 2021 ACM conference on fairness, accountability, and transparency*, pp. 353–362, 2021.

Levente Kocsis and Csaba Szepesvári. Bandit based monte-carlo planning. In *European conference on machine learning*, pp. 282–293. Springer, 2006.

Sanghack Lee and Elias Bareinboim. Structural causal bandits: Where to intervene? *Advances in neural information processing systems*, 31, 2018.

Lars Lorch, Jonas Rothfuss, Bernhard Schölkopf, and Andreas Krause. Dibs: Differentiable bayesian structure learning. *Advances in Neural Information Processing Systems*, 34:24111–24123, 2021.

Robert Northcott. Can anova measure causal strength? *The Quarterly Review of Biology*, 83(1): 47–55, 2008.

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

Judea Pearl and Dana Mackenzie. *The book of why: the new science of cause and effect*. Basic books, 2018.

Judea Pearl, Madelyn Glymour, and Nicholas P Jewell. *Causal inference in statistics: A primer*. John Wiley & Sons, 2016.

Jonas Peters, Dominik Janzing, and Bernhard Schölkopf. *Elements of causal inference: foundations and learning algorithms*. The MIT Press, 2017.

Ilan Price, Alvaro Sanchez-Gonzalez, Ferran Alet, Tom R Andersson, Andrew El-Kadi, Dominic Masters, Timo Ewalds, Jacklynn Stott, Shakir Mohamed, Peter Battaglia, et al. Probabilistic weather forecasting with machine learning. *Nature*, 637(8044):84–90, 2025.

Tian Qin, Tian-Zuo Wang, and Zhi-Hua Zhou. Rehearsal learning for avoiding undesired future. In *Advances in Neural Information Processing Systems*, volume 36, pp. 80517–80542, 2023.

Tian Qin, Tian-Zuo Wang, and Zhi-Hua Zhou. Gradient-based nonlinear rehearsal learning with multivariate alterations. In *AAAI conference on artificial intelligence*, volume 39, 2025.

Paul R Rosenbaum and Donald B Rubin. The central role of the propensity score in observational studies for causal effects. *Biometrika*, 70(1):41–55, 1983.

Stuart J Russell and Peter Norvig. *Artificial intelligence: a modern approach*. Pearson, 4 edition, 2020.

Amit Sharma and Emre Kiciman. Dowhy: An end-to-end library for causal inference. *arXiv preprint arXiv:2011.04216*, 2020.

Peter Spirtes, Clark N Glymour, and Richard Scheines. *Causation, prediction, and search*. MIT press, 2000.

Richard S Sutton and Andrew G Barto. *Reinforcement learning: An introduction*. MIT press, 2018.

Stratis Tsirtsis, Abir De, and Manuel Rodriguez. Counterfactual explanations in sequential decision making under uncertainty. *Advances in Neural Information Processing Systems*, 34:30127–30139, 2021.

Thomas S Verma and Judea Pearl. Equivalence and synthesis of causal models. In *Proceedings of the 6th Annual Conference on Uncertainty in Artificial Intelligence*, pp. 255–270, 1991.

Junzhe Zhang and Elias Bareinboim. Near-optimal reinforcement learning in dynamic treatment regimes. *Advances in Neural Information Processing Systems*, 32, 2019.

Xun Zheng, Bryon Aragam, Pradeep K Ravikumar, and Eric P Xing. Dags with no tears: Continuous optimization for structure learning. *Advances in neural information processing systems*, 31, 2018.

Zhi-Hua Zhou. Rehearsal: learning from prediction to decision. *Frontiers of Computer Science*, 16 (4):164352, 2022.

Zhi-Hua Zhou. Rehearsal: Learning from prediction to decision. Keynote at the CCF Conference on AI, 2023.

A DETAILED SETTINGS

Our experiments are conducted using Intel Xeon E-2288G CPUs, featuring 8 cores and 16 threads with a frequency of 3.7 GHz. The implementation is based on DOWHY (Sharma & Kiciman, 2020; Blöbaum et al., 2024). The code to reproduce our results will be made publicly available.

A.1 THE SYNTHETIC TASKS

The underlying SCM for the TRADER task governs the sequence of variables (V_1, V_2, V_3, V_4, Y) through the following structural equations:

$$\begin{aligned}
V_1 &\coloneqq N_1, \\
V_2 &\coloneqq V_1 \cdot N_2 + (1 - V_1) \cdot (1 - N_2), \\
V_3 &\coloneqq V_2 \cdot N_3 + (1 - V_2) \cdot (1 - N_3), \\
V_4 &\coloneqq V_3 \cdot N_4 + (1 - V_3) \cdot (1 - N_4), \\
Y &\coloneqq V_4 \cdot (1 - V_1) + (1 - V_4) \cdot N_Y,
\end{aligned}$$

where the noise terms follow $N_1 \sim \text{Bern}(\rho_1)$, $N_2 \sim \text{Bern}(\rho_2)$, $N_3 \sim \text{Bern}(\rho_3)$, $N_4 \sim \text{Bern}(\rho_4)$, and $N_Y \sim \text{Bern}(\rho_Y)$. To facilitate diversity across experimental trials, the task parameters are independently and uniformly sampled from the following intervals: $\rho_1 \in [0.4, 0.6]$, $\rho_i \in [0.7, 0.9]$ for $i \in \{2, 3, 4\}$, and $\rho_Y \in [0.3, 0.5]$. The variables V_2, V_3, and V_4 are actionable with feasible domains $\Delta_{V_2} = \Delta_{V_3} = \Delta_{V_4} = \{0, 1\}$. The variable V_1 is not observed at the time of decision. The desired domain is specified as $\mathcal{S} = \{1\}$. Here, V_1 represents the economic climate, and the chain $V_2 \to V_3 \to V_4$ models the progression from consumer demand to the final marketing strategy. The target Y denotes quarterly profit, whose structural equation explicitly encodes the interaction between strategy (V_4) and environment (V_1), implying that the profitability of a specific strategy relates to the prevailing economic state. The objective is to determine the actionable variables to maximize profit.

The underlying SCM for the FARMER task governs the sequence of variables (V_1, V_2, V_3, V_4, Y) thorough the following structural equations:

$$\begin{aligned}
V_1 &\coloneqq N_1, \\
V_2 &\coloneqq (1 - V_1) \cdot N_2, \\
V_3 &\coloneqq (1 - V_2) \cdot N_3, \\
V_4 &\coloneqq (1 - V_3) \cdot N_4, \\
Y &\coloneqq V_1 \cdot V_4 \cdot N_Y,
\end{aligned}$$

where the noise terms follow $N_1 \sim \text{Bern}(\beta_1)$, $N_2 \sim \text{Bern}(\beta_2)$, $N_3 \sim \text{Bern}(\beta_3)$, $N_4 \sim \text{Bern}(\beta_4)$, and $N_Y \sim \text{Bern}(\beta_Y)$. To facilitate diversity across experimental trials, the task parameters are independently and uniformly sampled from the following intervals: $\beta_i \in [0.60, 0.95]$ for $i \in \{1, 2, 3, 4\}$, and $\beta_Y \in [0.60, 0.95]$. The variables V_2, V_3, and V_4 are actionable with feasible domains $\Delta_{V_2} = \Delta_{V_3} = \Delta_{V_4} = \{0, 1\}$. The variable V_1 is not observed at the time of decision. The desired domain is specified as $\mathcal{S} = \{1\}$. In this context, V_1 represents sunlight exposure, while the chain $V_2 \to V_3 \to V_4$ models the natural water cycle affecting the soil. Specifically, intense sunlight (V_1) naturally reduces precipitation (V_2), which in turn increases evaporation (V_3), ultimately leading to low soil moisture (V_4). The target Y denotes crop yield. The structural equation for Y explicitly encodes the essential interaction between light (V_1) and water (V_4), implying that high productivity requires the simultaneous presence of both sunlight and adequate soil moisture. The objective is to determine the actionable variables to maximize crop yield.

The underlying SCM for the DOCTOR task governs the sequence of variables (V_1, V_2, V_3, V_4, Y) thorough the following structural equations:

$$\begin{aligned}
V_1 &\coloneqq N_1, \\
V_2 &\coloneqq N_2, \\
V_3 &\coloneqq V_1 \cdot V_2 \cdot (1 - N_3), \\
V_4 &\coloneqq N_4, \\
Y &\coloneqq V_4 \cdot (1 - V_1) + (1 - V_4) \cdot N_Y,
\end{aligned}$$

where the noise terms follow $N_1 \sim \text{Bern}(\gamma_1)$, $N_2 \sim \text{Bern}(\gamma_2)$, $N_3 \sim \text{Bern}(\gamma_3)$, $N_4 \sim \text{Bern}(\gamma_4)$, and $N_Y \sim \text{Bern}(\gamma_Y)$. To facilitate diversity across experimental trials, the task parameters are independently and uniformly sampled from the following intervals: $\gamma_1 \in [0.4, 0.6]$, $\gamma_i \in [0.1, 0.3]$ for $i \in \{2, 3, 4\}$, and $\gamma_Y \in [0.3, 0.5]$. The variables V_2, V_3, and V_4 are actionable with feasible domains $\Delta_{V_2} = \Delta_{V_3} = \Delta_{V_4} = \{0, 1\}$. The variable V_1 is not observed at the time of decision. The desired domain is specified as $\mathcal{S} = \{1\}$. In this context, V_1 represents the drug intolerance (or allergy gene), while V_2 represents an environmental trigger. V_3 denotes a symptom (e.g., a rash), which serves as a diagnostic indicator. The structural equation for V_3 implies that the symptom manifests primarily when both the intolerance (V_1) and the trigger (V_2) are present. V_4 represents the administration of a potent drug. The target Y denotes patient recovery. The equation for Y captures a critical medical contraindication: the drug (V_4) is effective for the general population ($V_1 = 0$) but is harmful or fatal to patients with the intolerance ($V_1 = 1$). The objective is to determine the actionable variables to maximize patient recovery.

A.2 The Bermuda Task

The Bermuda case study is derived from a real-world scenario involving the management of net coral ecosystem calcification in Bermuda, where environmental variables are recorded (Aglietti et al., 2020). The sequence of variables in this task are listed as follows:

- $Light$: bottom light levels;
- Tem: bottom temperature;
- Sal: sea surface salinity;
- DIC: seawater dissolved inorganic carbon;
- TA: seawater total alkalinity;
- Ω_A: seawater saturation with respect to aragonite;
- Nut: PC1 of NH_4, NiO_2 + NiO_3, SiO_4;
- $Chl\alpha$: sea surface chlorophyll-a;
- pH_{sw}: seawater pH;
- P_{CO_2}: seawater P_{CO_2};
- NEC: net ecosystem calcification.

The causal graph governing these variables is adopted from Courtney et al. (2017) and is illustrated in Figure 4. Consistent with previous studies Aglietti et al. (2020); Qin et al. (2023), the structural equations were obtained by performing linear regression on the 50 observations provided by Andersson & Bates (2018), and there are five actionable variables including DIC, TA, Ω_A, Nut, and

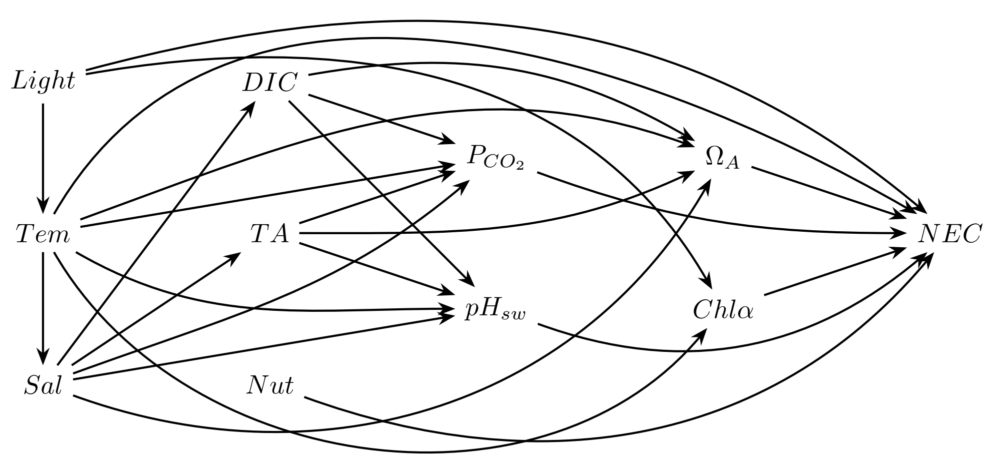

Figure 4: The causal graph for the Bermuda task.

$Chl\alpha$. Other variables are not observed at the time of decision. To ensure compatibility with our method, we discretize the continuous variables by dividing their value ranges into six equal-width bins. Accordingly, the feasible domains for alterations are set to $\{-2.5, -1.5, -0.5, 0.5, 1.5, 2.5\}$. The desired domain is specified as $\mathcal{S} = [1, 2]$. The objective is to determine which variables to alter in order to achieve a net ecosystem calcification (NEC) within the desired range, thereby promoting coral reef health and resilience.

B PROOF OF THEOREM 1

Theorem 1. *Let X and Y be two endogenous variables in an SCM. The following statements hold:*

1. $X \in \mathrm{Anc}(Y) \not\Rightarrow \dot{p}(X,Y) \neq 0$ *and* $\dot{p}(X,Y) \neq 0 \not\Rightarrow X \in \mathrm{Anc}(Y)$*;*
2. $\tau(X,Y) \neq 0 \not\Rightarrow \dot{p}(X,Y) \neq 0$ *and* $\dot{p}(X,Y) \neq 0 \not\Rightarrow \tau(X,Y) \neq 0$*;*
3. $\tau(X,Y) \neq 0 \not\Rightarrow \dot{p}(X,Y) \geq 0$ *and* $X \in \mathrm{Anc}(Y) \not\Rightarrow \dot{p}(X,Y) \geq 0$*.*

Proof. We prove each statement separately by constructing a counterexample.

Statement (a): $X \in \mathrm{Anc}(Y) \not\Rightarrow \dot{p}(X,Y) \neq 0$.

To show that $X \in \mathrm{Anc}(Y)$ does not imply $\dot{p}(X,Y) \neq 0$, it suffices to provide a case where a variable is an ancestor of another, yet its influence power on the latter is zero.

Consider the following SCM over the sequence of variables (V_1, Y):

$$\begin{aligned} V_1 &:= N_1, \\ Y &:= V_1 \cdot N_Y + (1 - V_1) \cdot (1 - N_Y), \end{aligned}$$

where $N_1 \sim \mathrm{Bern}(0.5)$, $N_Y \sim \mathrm{Bern}(0.5)$, V_1 is actionable with $\Delta_{V_1} = \{0, 1\}$, and the desired domain for Y is $\mathcal{S} = \{1\}$.

In the SCM, we have

$$V_1 \in \mathrm{Anc}(Y),$$

and

$$\begin{aligned} \dot{p}(V_1, Y) &= \max\nolimits_{v_1 \in \Delta_{V_1}} \mathcal{P}(Y = 1 | V_1 :- v_1) - \mathbb{E}_{v_1 \sim P(V_1)} \mathcal{P}(Y = 1 | V_1 = v_1) \\ &= \max\nolimits_{v_1 \in \Delta_{V_1}} P(Y = 1 | V_1 :- v_1) - P(Y = 1) \\ &= \max\{0.5, 0.5\} - 0.5 \\ &= 0. \end{aligned}$$

Thus, an ancestral relationship in the SCM does not imply non-zero influence power.

Statement (b): $\tau(X,Y) \neq 0 \not\Rightarrow \dot{p}(X,Y) \neq 0$.

To show that $\tau(X,Y) \neq 0$ does not imply $\dot{p}(X,Y) \neq 0$, it suffices to provide a case where a variable has a non-zero average causal effect on another, yet its influence power on the latter is zero.

Consider the following SCM over the sequence of variables (V_1, V_2, Y):

$$\begin{aligned} V_1 &:= N_1, \\ V_2 &:= V_1 \cdot N_2 + (1 - V_1) \cdot (1 - N_2), \\ Y &:= V_2 \cdot N_Y + (1 - V_2) \cdot (1 - N_Y), \end{aligned}$$

where $N_1, N_2, N_Y \overset{iid}{\sim} \mathrm{Bern}(0.9)$, V_1 and V_2 are actionable with $\Delta_{V_1} = \Delta_{V_2} = \{0, 1\}$, and the desired domain for Y is $\mathcal{S} = \{1\}$. This SCM corresponds to Example 1 in the main text.

In the SCM, we have

$$\begin{aligned} \tau(V_1, Y) &= P(Y = 1 | V_1 :- 1) - P(Y = 1 | V_1 :- 0) \\ &= 0.82 - 0.18 \\ &= 0.64, \end{aligned}$$

and

$$\begin{aligned}\dot{p}(V_1, Y) &= \max_{v_1 \in \Delta_{V_1}} \mathcal{P}(Y=1|V_1 :- v_1) - \mathbb{E}_{v_1 \sim P(V_1)} \mathcal{P}(Y=1|V_1 = v_1) \\ &= \max_{v_1 \in \Delta_{V_1}} \max_{v_2 \in \Delta_{V_2}} P(Y=1|V_2 :- v_2, V_1 :- v_1) \\ &\quad - \mathbb{E}_{v_1 \sim P(V_1)} \max_{v_2 \in \Delta_{V_2}} P(Y=1|V_2 :- v_2, V_1 = v_1) \\ &= \max\{0.9, 0.9\} - 0.9 \\ &= 0.\end{aligned}$$

Thus, a non-zero average causal effect in the SCM does not imply non-zero influence power. We also note that *Statement (b)* implies *Statement (a)*, as $\tau(X, Y) \neq 0$ implies $X \in \mathrm{Anc}(Y)$.

Statement (c): $\dot{p}(X, Y) \neq 0 \implies\!\!\!\!\!/\;\; \tau(X, Y) \neq 0$.

To show that $\dot{p}(X, Y) \neq 0$ does not imply $\tau(X, Y) \neq 0$, it suffices to provide a case where a variable has non-zero influence power on another, yet its average causal effect on the latter is zero.

Consider the following SCM over the sequence of variables (V_1, V_2, Y):

$$\begin{aligned}V_1 &\coloneqq N_1, \\ V_2 &\coloneqq (1 - V_1) \cdot N_2, \\ Y &\coloneqq V_1 \cdot V_2 \cdot N_Y,\end{aligned}$$

where $N_1, N_2, N_Y \overset{iid}{\sim} \mathrm{Bern}(0.5)$, V_1 and V_2 are actionable with $\Delta_{V_1} = \Delta_{V_2} = \{0, 1\}$, and the desired domain for Y is $\mathcal{S} = \{1\}$. This SCM corresponds to Example 3 in the main text.

In the SCM, we have

$$\begin{aligned}\tau(V_1, Y) &= P(Y=1|V_1 :- 1) - P(Y=1|V_1 :- 0) \\ &= 0.82 - 0.18 \\ &= 0.64,\end{aligned}$$

and

$$\begin{aligned}\dot{p}(V_1, Y) &= \max_{v_1 \in \Delta_{V_1}} \mathcal{P}(Y=1|V_1 :- v_1) - \mathbb{E}_{v_1 \sim P(V_1)} \mathcal{P}(Y=1|V_1 = v_1) \\ &= \max_{v_1 \in \Delta_{V_1}} \max_{v_2 \in \Delta_{V_2}} P(Y=1|V_2 :- v_2, V_1 :- v_1) \\ &\quad - \mathbb{E}_{v_1 \sim P(V_1)} \max_{v_2 \in \Delta_{V_2}} P(Y=1|V_2 :- v_2, V_1 = v_1) \\ &= \max\{0, 0.5\} - 0.25 \\ &= 0.25.\end{aligned}$$

Thus, non-zero influence power does not imply a non-zero average causal effect in the SCM.

Statement (d): $\dot{p}(X, Y) \neq 0 \implies\!\!\!\!\!/\;\; X \in \mathrm{Anc}(Y)$.

To show that $\dot{p}(X, Y) \neq 0$ does not imply $X \in \mathrm{Anc}(Y)$, it suffices to provide a case where a variable has non-zero influence power on another, yet it is not an ancestor of the latter.

Consider the following SCM over the sequence of variables (V_1, V_2, V_3, V_4, Y):

$$\begin{aligned}V_1 &\coloneqq N_1, \\ V_2 &\coloneqq N_2, \\ V_3 &\coloneqq V_1 \cdot V_2 \cdot (1 - N_3), \\ V_4 &\coloneqq N_4, \\ Y &\coloneqq V_4 \cdot (1 - V_1) + (1 - V_4) \cdot N_Y,\end{aligned}$$

where $N_1 \sim \mathrm{Bern}(0.5)$, $N_2, N_3, N_4 \overset{iid}{\sim} \mathrm{Bern}(0.1)$, $N_Y \sim \mathrm{Bern}(0.4)$, V_2, V_3, and V_4 are actionable with $\Delta_{V_2} = \Delta_{V_3} = \Delta_{V_4} = \{0, 1\}$, and the desired domain for Y is $\mathcal{S} = \{1\}$. This SCM corresponds to Example 2 in the main text.

In the SCM, we have

$$V_2 \notin \mathrm{Anc}(Y),$$

and

$$\begin{aligned}
\dot{p}(V_2, Y) &= \max_{v_2 \in \Delta_{V_2}} \mathcal{P}(Y=1|V_2 :- v_2) - \mathbb{E}_{v_2 \sim P(V_2)} \mathcal{P}(Y=1|V_2 = v_2) \\
&= \max_{v_2 \in \Delta_{V_2}} \max \Big\{ \max_{v_3 \in \Delta_{V_3}} \mathcal{P}(Y=1|V_3 :- v_3, V_2 :- v_2), \\
&\qquad \mathbb{E}_{v_3 \sim P(V_3|V_2 :- v_2)} \mathcal{P}(Y=1|V_3 = v_3, V_2 :- v_2) \Big\} \\
&\quad - \mathbb{E}_{v_2 \sim P(V_2)} \max \Big\{ \max_{v_3 \in \Delta_{V_3}} \mathcal{P}(Y=1|V_3 :- v_3, V_2 = v_2), \\
&\qquad \mathbb{E}_{v_3 \sim P(V_3|V_2 = v_2)} \mathcal{P}(Y=1|V_3 = v_3, V_2 = v_2) \Big\} \\
&= \max_{v_2 \in \Delta_{V_2}} \mathbb{E}_{v_3 \sim P(V_3|V_2 :- v_2)} \mathcal{P}(Y=1|V_3 = v_3, V_2 :- v_2) \\
&\quad - \mathbb{E}_{v_2 \sim P(V_2)} \mathbb{E}_{v_3 \sim P(V_3|V_2 = v_2)} \mathcal{P}(Y=1|V_3 = v_3, V_2 = v_2) \\
&= \max_{v_2 \in \Delta_{V_2}} \mathbb{E}_{v_3 \sim P(V_3|V_2 :- v_2)} \max_{v_4 \in \Delta_{V_4}} P(Y=1|V_4 :- v_4, V_3 = v_3, V_2 :- v_2) \\
&\quad - \mathbb{E}_{v_2 \sim P(V_2)} \mathbb{E}_{v_3 \sim P(V_3|V_2 = v_2)} \max_{v_4 \in \Delta_{V_4}} P(Y=1|V_4 :- v_4, V_3 = v_3, V_2 = v_2) \\
&= 0.68 - 0.518 \\
&= 0.162.
\end{aligned}$$

Thus, non-zero influence power does not imply an ancestral relationship in the SCM.
We also note that *Statement (d)* implies *Statement (c)*, as $X \notin \mathrm{Anc}(Y)$ implies $\tau(X, Y) = 0$.

Statement (e): $X \in \mathrm{Anc}(Y) \;\not\Longrightarrow\; \dot{p}(X, Y) \geq 0$.

To show that $X \in \mathrm{Anc}(Y)$ does not imply $\dot{p}(X, Y) \geq 0$, it suffices to provide a case where a variable is an ancestor of another, yet its influence power on the latter is negative.

Consider the following SCM over the sequence of variables (V_1, V_2, Y):

$$\begin{aligned}
V_1 &:= N_1, \\
V_2 &:= (1 - V_1) \cdot N_2, \\
Y &:= (V_1 \oplus V_2) \cdot N_Y,
\end{aligned}$$

where $N_1 \sim \mathrm{Bern}(0.5)$, $N_2, N_Y \overset{iid}{\sim} \mathrm{Bern}(0.8)$, V_2 is actionable with $\Delta_{V_2} = \{0, 1\}$, and the desired domain for Y is $\mathcal{S} = \{1\}$.

In the SCM, we have

$$V_2 \in \mathrm{Anc}(Y),$$

and

$$\begin{aligned}
\dot{p}(V_2, Y) &= \max_{v_2 \in \Delta_{V_2}} \mathcal{P}(Y=1|V_2 :- v_2) - \mathbb{E}_{v_2 \sim P(V_2)} \mathcal{P}(Y=1|V_2 = v_2) \\
&= \max_{v_2 \in \Delta_{V_2}} P(Y=1|V_2 :- v_2) - P(Y=1) \\
&= \max\{0.4, 0.4\} - 0.72 \\
&= -0.32.
\end{aligned}$$

Thus, an ancestral relationship in the SCM does not imply non-negative influence power.

Statement (f): $\tau(X, Y) \neq 0 \;\not\Longrightarrow\; \dot{p}(X, Y) \geq 0$.

To show that $\tau(X, Y) \neq 0$ does not imply $\dot{p}(X, Y) \geq 0$, it suffices to provide a case where a variable has a non-zero average causal effect on another, yet its influence power on the latter is negative.

Consider the following SCM over the sequence of variables (V_1, V_2, V_3, Y):

$$\begin{aligned}
V_1 &:= N_1, \\
V_2 &:= V_1 \cdot N_2 + (1 - V_1) \cdot (1 - N_2), \\
V_3 &:= V_2 \cdot N_3 + (1 - V_2) \cdot (1 - N_3), \\
Y &:= V_3 \cdot (1 - V_1) + (1 - V_3) \cdot N_Y,
\end{aligned}$$

where $N_1 \sim \text{Bern}(0.5)$, $N_2, N_3 \overset{iid}{\sim} \text{Bern}(0.9)$, $N_Y \sim \text{Bern}(0.4)$, V_2 and V_3 are actionable with $\Delta_{V_2} = \{0, 1\}$ and $\Delta_{V_3} = \{0, 1\}$, and the desired domain for Y is $\mathcal{S} = \{1\}$. This SCM corresponds to Example 4 in the main text.

In the SCM, we have

$$\begin{aligned}\tau(V_2, Y) &= P(Y = 1|V_2 :- 1) - P(Y = 1|V_2 :- 0)\\ &= 0.49 - 0.41\\ &= 0.08,\end{aligned}$$

and

$$\begin{aligned}\dot{p}(V_2, Y) &= \max_{v_2 \in \Delta_{V_2}} \mathcal{P}(Y = 1|V_2 :- v_2) - \mathbb{E}_{v_2 \sim P(V_2)} \mathcal{P}(Y = 1|V_2 = v_2)\\ &= \max_{v_2 \in \Delta_{V_2}} \max_{v_3 \in \Delta_{V_3}} P(Y = 1|V_3 :- v_3, V_2 :- v_2)\\ &\quad - \mathbb{E}_{v_2 \sim P(V_2)} \max_{v_3 \in \Delta_{V_3}} P(Y = 1|V_3 :- v_3, V_2 = v_2)\\ &= \max\{0.5, 0.5\} - 0.65\\ &= -0.15.\end{aligned}$$

Thus, a non-zero average causal effect in the SCM does not imply non-negative influence power. We also note that *Statement (f)* implies *Statement (e)*, as $\tau(X, Y) \neq 0$ implies $X \in \text{Anc}(Y)$. □

C Proof of Proposition 1

Proposition 1. *Assume causal sufficiency, i.e., the joint distribution $P(\mathbf{V}, Y)$ is induced by an acyclic SCM $\mathcal{M}$ with mutually independent background noises, and positivity, i.e., $P(V_i|\text{PA}_i) > 0$ in the support of P, $\forall 1 \leq i \leq d$. Then, the expression in Equation* (11) *is consistent to the joint probability dictated by the SCM $\mathcal{M}_{\mathbf{A}}$ where variables $\mathbf{A}$ are altered. Furthermore, the expression in Equation* (12) *is consistent to the AUF probability dictated by the SCM $\mathcal{M}_{\mathbf{A}}$ where variables $\mathbf{A}$ are altered and variables $\mathbf{O}$ are observed.*

Proof. Recall from Equation (11), the joint distribution conditioned on the alteration set $\hat{\mathbf{A}}$ is expressed as $P(\mathbf{X}|\hat{\mathbf{A}}) = \prod_{X_i \in \mathbf{A}} \delta(X_i) \prod_{X_i \in \mathbf{X} \setminus \mathbf{A}} P(X_i|X_1, \ldots, X_{i-1})$. As the sequence is topologically consistent with the underlying SCM, and the SCM is assumed to be acyclic, the value of each variable X_i depends solely on its parents PA_i. Consequently, $P(X_i|X_1, \ldots, X_{i-1}) = P(X_i|\text{PA}_i)$. Substituting this back into the product shows that $P(\mathbf{X}|\hat{\mathbf{A}}) = \prod_{X_i \in \mathbf{A}} \delta(X_i) \prod_{X_i \in \mathbf{X} \setminus \mathbf{A}} P(X_i|\text{PA}_i)$. By invoking the manipulation theorem (i.e., Theorem 3.6 in Spirtes et al. (2000)), we have that $P(\mathbf{X}|\hat{\mathbf{A}})$ is exactly the probability of $\mathbf{X}$ under alteration of $\mathbf{A}$. Moreover, the quantity $P(Y \in \mathcal{S}|\hat{\mathbf{A}}, \mathbf{O})$ in Equation (12) is fully determined by $P(\mathbf{X}|\hat{\mathbf{A}})$, and therefore Equation (12) indeed gives to the true AUF probability dictated by the underlying SCM. □

D Additional Related Work

The rehearsal paradigm was introduced by Zhou (2022), building on the concept of influence (Zhou, 2023), This paradigm advocates for mentally simulating future possibilities in order to find alterations that positively influence the future target before making a final decision. This is analogous to how human cognitive process prepares for future events (Driskell et al., 1994). Motivated by this, Qin et al. (2023) proposed the first rehearsal learning approach, wherein the restriction of directionality is relaxed and *structural rehearsal models* capable of accommodating bi-directional interactions are developed. Several subsequent studies have addressed issues such as non-stationarity and non-linearity in rehearsal learning (Du et al., 2024; Qin et al., 2025), requiring that the structure of the underlying equations are provided by experts. Besides, while the forward-looking decision-making problem is also conceptually related to markov decision processes in reinforcement learning (Sutton & Barto, 2018), a key distinction is that the AUF problem operates under a "no going back" constraint. Unlike in many RL settings where an agent can revisit states, the past variables cannot be changed in our context. Our approximation method is particularly inspired by Monte Carlo Tree Search (MCTS) (Browne et al., 2012), which excel at planning in large state spaces by simulating future trajectories, making them well-suited for the challenges of the AUF problem.