

Backtracking Counterfactuals

Julius von Kügelgen

*Max Planck Institute for Intelligent Systems Tübingen, Germany
University of Cambridge*

JVK@TUEBINGEN.MPG.DE

Abdirisak Mohamed

*University of Maryland
SAP Labs, LLC*

AMOHAM70@UMD.EDU

Sander Beckers

University of Tübingen

SREKCEBREDNAS@GMAIL.COM

Editors: Mihaela van der Schaar, Dominik Janzing and Cheng Zhang

Abstract

Counterfactual reasoning—envisioning hypothetical scenarios, or possible worlds, where some circumstances are different from what (f)actually occurred (counter-to-fact)—is ubiquitous in human cognition. Conventionally, counterfactually-altered circumstances have been treated as “small miracles” that locally violate the laws of nature while sharing the same initial conditions. In Pearl’s structural causal model (SCM) framework this is made mathematically rigorous via *interventions* that modify the causal laws while the values of exogenous variables are shared. In recent years, however, this purely interventionist account of counterfactuals has increasingly come under scrutiny from both philosophers and psychologists. Instead, they suggest a *backtracking* account of counterfactuals, according to which *the causal laws remain unchanged* in the counterfactual world; differences to the factual world are instead “backtracked” to altered initial conditions (exogenous variables). In the present work, we explore and formalise this alternative mode of counterfactual reasoning within the SCM framework. Despite ample evidence that humans backtrack, the present work constitutes, to the best of our knowledge, the first general account and algorithmisation of backtracking counterfactuals. We discuss our backtracking semantics in the context of related literature and draw connections to recent developments in explainable artificial intelligence (XAI).

Keywords: Causal reasoning, backtracking, counterfactual explanations, explainable AI, XAI

1. Introduction

In a deterministic world, everything that happens is uniquely determined by the laws of nature and the initial (or background) conditions.¹ Counterfactuals invite us to imagine what the world would look like if some events which did occur, had in fact not occurred. As a result, in order to make sense of counterfactuals in a deterministic world one is immediately presented with the following dilemma: for events to have been different, either the laws of nature would have had to be violated, or the background conditions would have had to be different.

Following in the footsteps of Lewis (1973, 1979), the majority of philosophers have opted for the first option: counterfactuals are to be evaluated by imagining “small miracles” that ensure those events which are counter-to-fact to occur by locally violating the laws of nature, thereby disconnecting these events from their causes, and keeping the background conditions fixed. In other words, counterfactuals are *non-backtracking*. Pearl (2009, p.239) objected to Lewis’s miracles and the possible-world semantics that comes with it, instead replacing these with structural equations and the well-behaved notion of an intervention that those allow:

1. Even if the world is fundamentally non-deterministic at the quantum level, the point still holds for deterministic *models* of the world. Causal models in particular can be given a deterministic interpretation, see (2.2).

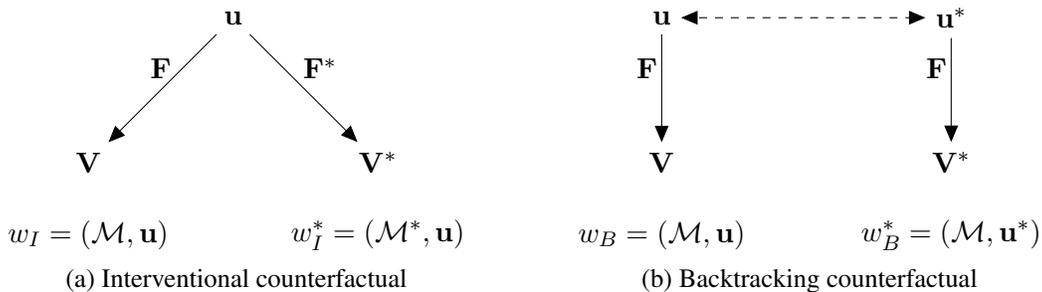


Figure 1: **Illustration of the main conceptual difference between interventional and backtracking counterfactuals.** (a) For *interventional* counterfactuals, the factual world w_I and counterfactual world w_I^* share the exact same background conditions \mathbf{u} . Potential contradictions between the factual outcome \mathbf{V} and the counterfactual outcome \mathbf{V}^* are resolved through changes to the causal laws \mathbf{F} (by means of intervention), giving rise to the modified laws \mathbf{F}^* and submodel \mathcal{M}^* . (b) For *backtracking* counterfactuals, on the other hand, the factual world w_B and counterfactual world w_B^* share the same unmodified causal laws \mathbf{F} . To allow for the factual and counterfactual outcomes \mathbf{V} and \mathbf{V}^* to differ, the respective background conditions \mathbf{u} and \mathbf{u}^* may differ, too, with a preference for minimal changes, that is, for the two worlds to be as close as possible.

Lewis’s elusive “miracles” are replaced by principled minisurgeries, [...] which represent the minimal change (to a model) necessary for establishing the antecedent.

Still, Pearl’s *interventional* counterfactuals are also a commitment to the first option of our dilemma.

Recently this status quo has come under pressure by several philosophers who argue in favour of the second option (Dorr, 2016; Loewer, 2007, 2020; Esfeld, 2021): counter-to-fact events are to be explained by imagining small changes to the background conditions that would result in the counterfactual events instead of the actual ones, while the laws of nature remain unchanged. In other words, counterfactuals are *backtracking*.² Others propose a combination of both options, suggesting that each is appropriate under different circumstances (Fisher, 2017a,b; Woodward, 2021).

This dispute is not merely academic: empirical work by psychologists now confirms that, depending on the context, people indeed switch between interpreting counterfactuals according to the first and the second option (Rips, 2010; Gerstenberg et al., 2013; Lucas and Kemp, 2015). Historical explanations often rely on counterfactuals as well, and it has been argued that “counterfactuals in history are backtracking” (Reiss, 2009, p. 719). Moreover, counterfactual explanations are taking up a prominent role in the work on explainable AI; their interpretation thus has significant repercussions for real-life applications (Wachter et al., 2017).

While non-backtracking, interventional counterfactuals have been given a well-defined semantics within Pearl’s framework of structural causal models (SCMs), the same cannot be said for backtracking counterfactuals. In this paper we offer the first general formal semantics for backtracking counterfactuals within SCMs. Our general semantics can be made more specific in various ways, depending on the particular purpose that is at stake. We do not view these backtracking semantics as an alternative to the standard non-backtracking semantics: the value of the latter for causal reasoning is undeniable. Rather, we see both semantics as being fit for different reasoning tasks.

it is appropriate to use backtracking counterfactuals to answer [...] how the past would have had to have been different had the present been different. [...] backtracking counterfactuals are important in diagnostic reasoning. However, this does not mean that it is

2. or *observational*, we use both terms interchangeably; likewise for non-backtracking and interventional

misguided to use non-backtracking counterfactuals to answer other sorts of questions such as those having to do with whether Cs cause Es. The two kinds of counterfactuals are just different, with different truth conditions (Woodward, 2021, p. 206)

Structure and Contributions. We start by reviewing the SCM framework and its interventional semantics for counterfactual reasoning (§ 2). In § 3, we then provide a comprehensive account of the alternative, backtracking mode of counterfactual reasoning by explaining the high-level intuition (§ 3.1, see Fig. 1 for an overview), introducing a formal semantics (§ 3.2), discussing some of the design choices and desiderata (§ 3.3), and providing some first theoretical insights (§ 3.4). We compare our semantics to previous attempts and discuss related work from other fields in § 4. We then highlight the use of backtracking counterfactuals for explainable artificial intelligence (XAI) (§ 5), and conclude with an outlook and suggestions for future work (§ 6).

2. Preliminaries: Structural Causal Models and Interventional Counterfactuals

The following definitions of SCMs and their counterfactual semantics follow Pearl (2009, § 7.1).

A **causal model** is a triple $\mathcal{M} = (\mathbf{U}, \mathbf{V}, \mathbf{F})$ where: (i) \mathbf{U} is a set $\{U_1, \dots, U_m\}$ of **exogenous** (background) variables determined by factors outside the model; (ii) \mathbf{V} is a set $\{V_1, V_2, \dots, V_n\}$ of **endogenous** variables determined by variables in the model, that is, variables in $\mathbf{U} \cup \mathbf{V}$; and (iii) \mathbf{F} is a set of functions $\{f_1, f_2, \dots, f_n\}$ s.t. each f_i is a mapping from (the respective domains of) $\mathbf{U}_i \cup \mathbf{PA}_i$ to V_i , where $\mathbf{U}_i \subseteq \mathbf{U}$ and $\mathbf{PA}_i \subseteq \mathbf{V} \setminus \{V_i\}$, and the entire set \mathbf{F} forms a mapping from \mathbf{U} to \mathbf{V} . In other words, each f_i in the **structural equations**, or **causal laws**,

$$V_i := f_i(\mathbf{PA}_i, \mathbf{U}_i) \quad i = 1, \dots, n, \quad (2.1)$$

assigns a value to V_i that depends on (the values of) a select set of variables in $\mathbf{U} \cup \mathbf{V}$, and the entire set \mathbf{F} has a unique solution $\mathbf{V}(\mathbf{u})$. The latter is ensured, e.g., in acyclic (“recursive”) systems.

The **causal diagram** $G(\mathcal{M})$ associated with causal model \mathcal{M} is the directed graph in which each node corresponds to a variable and directed edges point from members of \mathbf{PA}_i and \mathbf{U}_i toward V_i . Since the exogenous (background) variables \mathbf{U} are typically unobserved, it is common to only consider the subset of $G(\mathcal{M})$ corresponding to its projection onto \mathbf{V} , where shared exogenous parents between some V_i and V_j are indicated with a bi-directed arrow $V_i \leftrightarrow V_j$.

For a subset of endogenous variables $\mathbf{X} \subseteq \mathbf{V}$ and a realisation \mathbf{x} thereof, the **submodel** $\mathcal{M}_{\mathbf{x}}$ of \mathcal{M} is the model $\mathcal{M}_{\mathbf{x}} = (\mathbf{U}, \mathbf{V}, \mathbf{F}_{\mathbf{x}})$ where $\mathbf{F}_{\mathbf{x}} = \{f_i : V_i \notin \mathbf{X}\} \cup \{\mathbf{X} := \mathbf{x}\}$. The **effect of action** $do(\mathbf{X} = \mathbf{x})$ on \mathcal{M} is given by the submodel $\mathcal{M}_{\mathbf{x}}$. A **causal world** w is a pair $(\mathcal{M}, \mathbf{u})$ where \mathcal{M} is a causal model and \mathbf{u} is a particular realization of the background variables \mathbf{U} . The **potential response** of $\mathbf{Y} \subseteq \mathbf{V}$ to action $do(\mathbf{X} = \mathbf{x})$ in world $w = (\mathcal{M}, \mathbf{u})$, denoted $\mathbf{Y}_{\mathbf{x}}(\mathbf{u})$, is the solution for \mathbf{Y} of the set of equations $\mathbf{F}_{\mathbf{x}}$, that is, $\mathbf{Y}_{\mathbf{x}}(\mathbf{u}) = \mathbf{Y}_{\mathcal{M}_{\mathbf{x}}}(\mathbf{u})$. The **counterfactual** sentence “ \mathbf{Y} would be \mathbf{y} (in situation \mathbf{u}), had \mathbf{X} been \mathbf{x} ” is then interpreted as the equality $\mathbf{Y}_{\mathbf{x}}(\mathbf{u}) = \mathbf{y}$. The part “had \mathbf{X} been \mathbf{x} ” is called the (counterfactual) **antecedent**.

A **probabilistic causal model** is a distribution over causal worlds, that is, a pair $(\mathcal{M}, P(\mathbf{U}))$ where \mathcal{M} is a causal model and $P(\mathbf{U})$ is a probability function defined over the domain of \mathbf{U} . The function $P(\mathbf{U})$, together with the fact that each endogenous variable is a function of \mathbf{U} , defines a **distribution over endogenous variables**: for any $\mathbf{Y} \subseteq \mathbf{V}$ we have³

$$P(\mathbf{y}) := P(\mathbf{Y} = \mathbf{y}) = \sum_{\mathbf{u}} P(\mathbf{u}) \mathbf{1}_{\{\mathbf{Y}(\mathbf{u})=\mathbf{y}\}} \quad (2.2)$$

3. To focus on the main points, ease the notation, and avoid measure theoretic details, we present all formulae for *discrete* random variables in terms of *sums* of probability *mass* functions; the analogue for *continuous* random variables would involve *integrals* of probability density functions (assuming existence of densities w.r.t. the Lebesgue measure).

where $\mathbf{1}$ denotes the indicator function. The **probability of counterfactuals** is defined analogously, through the potential responses induced by different submodels: for any not necessarily disjoint sets of variables $\mathbf{Y}, \mathbf{X}, \mathbf{Z}, \mathbf{W} \subseteq \mathbf{V}$, we have

$$P(\mathbf{Y}_{\mathbf{x}} = \mathbf{y}, \mathbf{Z}_{\mathbf{w}} = \mathbf{z}) = \sum_{\mathbf{u}} P(\mathbf{u}) \mathbf{1}_{\{\mathbf{Y}_{\mathbf{x}}(\mathbf{u})=\mathbf{y}\}} \mathbf{1}_{\{\mathbf{Z}_{\mathbf{w}}(\mathbf{u})=\mathbf{z}\}}. \quad (2.3)$$

In particular, $P(\mathbf{Y}_{\mathbf{x}} = \mathbf{y}, \mathbf{X} = \mathbf{x}')$ and $P(\mathbf{Y}_{\mathbf{x}} = \mathbf{y}, \mathbf{Y}_{\mathbf{x}'} = \mathbf{y}')$ are well defined using SCM semantics, even though $\mathbf{x} \neq \mathbf{x}'$ may be incompatible and thus cannot be measured simultaneously.⁴

To notationally distinguish the factual variables \mathbf{V} from counterfactual versions, or copies, thereof, the latter are sometimes also denoted with an asterisk \mathbf{V}^* (Balke and Pearl, 1994).

Of special interest are counterfactuals of the form $P(\mathbf{Y}_{\mathbf{x}^*} = \mathbf{y}^* \mid \mathbf{z})$ that are conditional on a (f)actual observation \mathbf{z} . These can be computed via the following three-step procedure:

1. **Abduction:** Update $P(\mathbf{U})$ by the evidence \mathbf{z} to obtain $P(\mathbf{U} \mid \mathbf{z})$.
2. **Action:** Modify \mathcal{M} by the action $do(\mathbf{x}^*)$ to obtain the submodel $\mathcal{M}_{\mathbf{x}^*}$.
3. **Prediction:** Use the modified model $(\mathcal{M}_{\mathbf{x}^*}, P(\mathbf{U} \mid \mathbf{z}))$ to compute the probability of \mathbf{Y}^* .

Example 1 (Interventional counterfactual) Consider a probabilistic causal model $(\mathcal{M}, P(\mathbf{U}))$ with $\mathbf{V} = \{X, Y, Z\}$, $\mathbf{U} = \{U_X, U_Y, U_Z\}$, causal laws \mathbf{F} given by

$$X := U_X, \quad Y := X + U_Y, \quad Z := X + Y + U_Z, \quad (2.4)$$

and $P(\mathbf{U})$ being the multivariate standard isotropic Gaussian distribution $\mathcal{N}(\mathbf{0}, \mathbf{I}_3)$.

Suppose that we make the factual observation $(X = 1, Y = 2, Z = 2)$ and wish to reason about the (interventional) counterfactual “what would have been, had Y been 3”, that is, we are interested in $P((X^*, Z^*)_{Y^*=3} \mid X = 1, Y = 2, Z = 2)$, see Fig. 2a for a visualisation.

To this end, we follow the aforementioned three-step procedure:

1. **Abduction:** we find the posterior $P(\mathbf{U} \mid X = 1, Y = 2, Z = 2)$ to be a point mass on $(U_X = 1, U_Y = 1, U_Z = -1)$.
2. **Action:** we modify \mathcal{M} to obtain the submodel $\mathcal{M}_{Y^*=3}$ as

$$X^* := U_X, \quad Y^* := 3, \quad Z^* := X^* + Y^* + U_Z; \quad (2.5)$$

Note that this step alters the mechanism for the counterfactual antecedent Y^* and removes its dependence on X^* and U_Y , as shown in Fig. 2a.

3. **Prediction:** we compute the push-forward of $P(\mathbf{U} \mid X = 1, Y = 2, Z = 2)$ via (2.5) which yields a point mass on $(X^* = 1, Z^* = 3)$. That is, X (being a parent of Y) would have remained unaffected, but Z (being a child of Y) would have increased by one.⁵

Remark 1 (Forward-tracking) In acyclic causal models, any manipulation $\mathbf{F}_{\mathbf{x}}$ of the causal laws \mathbf{F} such as $Y^* := 3$ in (2.5) only has “downstream” effects: the change only propagates to descendants of the intervened-upon variables, while any non-descendants remain unaffected. In this sense, interventions and thus also interventional counterfactuals are purely forward-tracking.

The *interventional* approach to computing counterfactuals in SCMs thus assumes that the background conditions \mathbf{u} are shared between the factual and counterfactual worlds and instead relies on modifying the causal laws (the *action* step) to explain possible discrepancies (such as the difference $Y = 2 \neq Y^* = 3$ in Ex. 1), as illustrated in Figs. 1a and 2a. As Pearl (2009, p.205) puts it:

4. Some therefore consider statements about such expressions as fundamentally unscientific, see, e.g., Dawid (2000).

5. For a comparison with the corresponding observational and interventional distributions, see Appx. A.

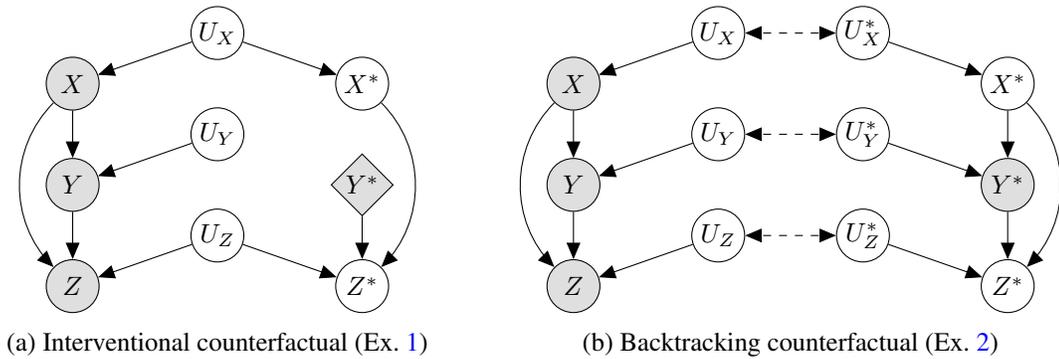


Figure 2: **Graphical comparison between the interventional and backtracking approaches to computing the counterfactual $P(X^*, Z^* | Y^*, X, Y, Z)$.** Shaded nodes are “observed” (or hypothesised to be known, in the case of counterfactual/starred variables), while white nodes are unobserved/latent. (a) For *interventional* counterfactuals, the counterfactual antecedent Y^* is assumed to arise from a small miracle in the form of a (local) change to the causal laws, illustrated by the diamond node. This intervention removes the dependence of Y^* on its causal parents X^* and U_Y , hence the missing arrows into Y^* . The exogenous variables are assumed to be shared, giving rise to a so-called “twin network” (Balke and Pearl, 1994). As a consequence, non-descendants of Y^* (here, only the parent X^*), remain the same and only descendants of Y^* (here, only the child Z^*) are affected. (b) For *backtracking* counterfactuals, the causal relationship between the counterfactual antecedent Y^* and its parents X^* and U_X^* is preserved. This necessitates the introduction of new counterfactual exogenous variables (U_X^* , U_Y^* , U_Z^*) which should be “close” to the original (U_X , U_Y , U_Z) but may differ from it to absorb, or explain, the difference between Y and Y^* . As a result of such backtracking, not only descendants (Z^*) but also ancestors (X^*) may change.

[It] interprets the counterfactual phrase “had \mathbf{X} been \mathbf{x} ” in terms of a *hypothetical modification of the equations in the model*; it simulates an *external action* (or spontaneous change) that modifies the actual course of history and enforces the condition “ $\mathbf{X} = \mathbf{x}$ ” with minimal *change of mechanisms*. This [...] permits \mathbf{x} to differ from the current value of $\mathbf{X}(\mathbf{u})$ without creating logical contradiction; it also suppresses abductive inferences (or *backtracking*) from the counterfactual antecedent $\mathbf{X} = \mathbf{x}$ [*emphasis ours*]

3. Backtracking Counterfactuals

We now explore and formalise an alternative, non-interventional mode of counterfactual reasoning that does not involve such “change of mechanism” by “external action”, and instead relies on so-called *backtracking* all changes to changes in the values of exogenous variables.

3.1. Intuition and Main Idea

The main idea behind backtracking counterfactuals—and its core conceptual difference to interventional counterfactuals—is that the causal laws, not the background conditions, are shared between the factual and counterfactual worlds. Consequently, the exogenous variables need to be allowed to differ to accommodate possibly contradictory facts and counterfacts. This is illustrated in Fig. 1b.

The name “backtracking” refers to the act of updating upstream variables to explain counterfactuals without breaking the causal laws. Such backtracking is not needed for interventional counterfactuals for which the modified equations take care of this. Let us first explore this idea through an example.

Example 2 (Backtracking counterfactual) Consider the same model (2.4) and factual observation ($X = 1, Y = 2, Z = 2$) from Ex. 1. Suppose that now we are instead interested in the backtracking counterfactual “what would have been, had Y instead been observed to be $Y^* = 3$ ”.

From abduction on the factual observation, we obtain the same posterior $P(\mathbf{U} \mid X = 1, Y = 2, Z = 2)$ being a point mass on $(U_X = 1, U_Y = 1, U_Z = -1)$. However, in contrast to (2.5), the structural equations in the counterfactual world now take the form:

$$X^* := U_X^*, \quad Y^* := X^* + U_Y^* = 3, \quad Z^* := X^* + Y^* + U_Z^*; \quad (3.1)$$

We highlight the following two key differences between (3.1) and its interventional counterpart (2.5): (i) the causal laws remain fully in place, as apparent from the dependence of Y^* on X^* and U_Y^* in (3.1); (ii) the exogenous variables are not (necessarily) shared, as apparent from the introduction of new counterfactual background variables (U_X^*, U_Y^*, U_Z^*) in (3.1).

Now, in order for $Y^* := X^* + U_Y^* = 3$ to hold, at least one of X^* or U_Y^* need to differ from their factual values of $X = U_Y = 1$. In other words, with the causal laws in the counterfactual world still being in place, the change from $Y = 2$ to $Y^* = 3$ needs to be explained by a change in X , U_Y , or both. In fact, there is an infinite number of such changed (X^*, U_Y^*) given by the solutions to $X^* + U_Y^* = 3$. Let us explore a few simple options:

- (i) We explain the change in Y entirely by a change in U_Y : this can be achieved by keeping X^* and hence also U_X^* equal to their factual values of one and setting $U_Y^* = 2$. Then U_Z^* could remain equal to its factual value of $U_Z = -1$, leading to $Z^* := X^* + Y^* + U_Z^* = 1 + 3 - 1 = 3$ which is identical to the interventional counterfactual computed in Ex. 1.⁶
- (ii) We explain the change in Y entirely by a change in X^* : we keep $U_Y^* = U_Y = 1$ and set $X^* = 2$. As a consequence, for the causal law $X^* := U_X^*$ to still be satisfied, we also need to adjust U_X and set $U_X^* = 2$. Again, we can keep $U_Z^* = U_Z = -1$, leading to $Z^* := X^* + Y^* + U_Z^* = 2 + 3 - 1 = 4$ which differs from the interventional counterfactual computed in Ex. 1!
- (iii) We explain the change in Y by a change in both X^* and U_Y^* : for this we set $U_Y^* = X^* = 1.5$, adjust U_X accordingly to $U_X^* = 1.5$, keep $U_Z^* = U_Z = -1$, and obtain $Z^* := X^* + Y^* + U_Z^* = 1.5 + 3 - 1 = 3.5$, which is yet another outcome for Z^* .

Before returning to Ex. 2 later, let us summarise some of the points it illustrates already:

First, to explain the change in $Y^* \neq Y$ while keeping the causal laws intact, we had to change at least one of the (endogenous or exogenous) parents of Y , that is, to propagate the change upstream in the causal hierarchy—this is exactly what *backtracking* refers to (cf. Remark 1).

Second, unlike for the interventional counterfactual in Ex. 1, the backtracking counterfactual in Ex. 2 does not lead to a unique solution. Since the background variables are allowed to differ, there can be multiple ways of setting $\mathbf{U}^* \neq \mathbf{U}$ that satisfy the causal laws and agree with the counterfactual antecedent. When changing only the exogenous variable U_Y^* corresponding to the counterfactual antecedent Y^* as in case (i), we reached the same conclusion as when interpreting such change in interventional terms as in Ex. 1. When changing an ancestor that influences a descendant ($X^* \rightarrow Z^*$) as in cases (ii) and (iii), however, we reached different conclusions from Ex. 1. This highlights that backtracking counterfactuals may or may not agree with their interventional counterparts. Note also that it was not necessary to change the exogenous variable U_Z^* of the descendant variable Z of Y from its factual value U_Z , even though this would in principle also be possible.

6. Except for the detail that, here, $Y^* = 2$ arises from a change in U_Y , instead of through an intervention $do(Y^* = 2)$ as in Ex. 1. However, the resulting implications for the endogenous variables are the same.

Recall that in the factual world we have $(U_X, U_Y, U_Z) = (1, 1, -1)$. We then constructed three counterfactual worlds that are consistent with $Y^* = 2$ and the causal laws:⁷

$$(U_X^*, U_Y^*, U_Z^*) = \begin{cases} (1, 2, -1) & \implies (X^*, Z^*) = (1, 3) & \text{for case (i),} \\ (2, 1, -1) & \implies (X^*, Z^*) = (2, 4) & \text{for case (ii),} \\ (1.5, 1.5, -1) & \implies (X^*, Z^*) = (1.5, 3.5) & \text{for case (iii).} \end{cases} \quad (3.2)$$

In order to pick one of these worlds, or form a weighted average of their predictions, we require a notion of preference between worlds, such as a closeness or similarity measure. This will be our starting point for formalising the semantics of backtracking counterfactuals.

3.2. Formal Semantics

In line with the probabilistic causal modelling framework, we now introduce the missing object for formal probabilistic reasoning about backtracking counterfactuals: an a priori similarity measure between factual and counterfactual worlds in the form of a *backtracking conditional* $P_B(\mathbf{U}^* \mid \mathbf{U})$. This is a collection of probability functions $P_B(\mathbf{U}^* \mid \mathbf{U} = \mathbf{u})$, one for each \mathbf{u} in the domain of \mathbf{U} , quantifying the likelihood of each counterfactual world \mathbf{u}^* given factual world \mathbf{u} ,⁸ prior to any factual (or counterfactual) observations. This provides us with a very general framework to encode various notions of cross-world similarity, depending on the context and the involved variables' domains. We remain agnostic to the exact choice of $P_B(\mathbf{U}^* \mid \mathbf{U})$ for now and defer further discussion of desiderata, properties, and specific choices thereof to § 3.3.

Together with the prior $P(\mathbf{U})$, the backtracking conditional $P_B(\mathbf{U}^* \mid \mathbf{U})$ induces a *joint distribution over factual and counterfactual worlds* given by:

$$P_B(\mathbf{U}^*, \mathbf{U}) = P(\mathbf{U})P_B(\mathbf{U}^* \mid \mathbf{U}) \quad (3.3)$$

We can now define the (joint) probability of backtracking counterfactuals similarly to (2.3).

Definition 2 (Probability of backtracking counterfactuals) *Let $(\mathcal{M}, P(\mathbf{U}))$ be a probabilistic causal model and $P_B(\mathbf{U}^* \mid \mathbf{U})$ a backtracking conditional, where \mathbf{U}^* denotes the counterfactual version of \mathbf{U} defined over the same domain. For any not necessarily disjoint subsets of variables $\mathbf{Y}, \mathbf{Z} \subseteq \mathbf{V}$ and any realisations \mathbf{y}^*, \mathbf{z} thereof, the probability of backtracking counterfactuals is given by:*

$$P_B(\mathbf{y}^*, \mathbf{z}) := P_B(\mathbf{Y}^* = \mathbf{y}^*, \mathbf{Z} = \mathbf{z}) = \sum_{(\mathbf{u}^*, \mathbf{u})} P_B(\mathbf{u}^*, \mathbf{u}) \mathbf{1}_{\{\mathbf{Y}^*(\mathbf{u}^*)=\mathbf{y}^*\}} \mathbf{1}_{\{\mathbf{Z}(\mathbf{u})=\mathbf{z}\}}. \quad (3.4)$$

Any other quantities of interest can then be derived from (3.4) by standard probabilistic inference (that is, through marginalisation and conditioning). In particular, we can now answer backtracking or observational counterfactuals of the form: “given that we factually observed \mathbf{Z} to be \mathbf{z} , what would be the probability that \mathbf{Y} would be \mathbf{y}^* , had we *observed* \mathbf{X} to be \mathbf{x}^* ?”. Provided that $P_B(\mathbf{x}^*, \mathbf{z}) > 0$, we obtain the corresponding expression $P_B(\mathbf{y}^* \mid \mathbf{x}^*, \mathbf{z})$ through the following three-step procedure, which loosely mirrors that for interventional counterfactuals given in § 2:

7. In fact, the (infinite) set of all valid counterfactual worlds is given by all $(U_X^*, 3 - U_X^*, U_Z^*)$ in the domain of \mathbf{U} , with corresponding predictions $(X^*, Z^*) = (U_X^*, 3 + U_X^* + U_Z^*)$.

8. Since the causal laws, and thus the causal model \mathcal{M} , remain unchanged for backtracking counterfactuals, we will also simply refer to instantiations of \mathbf{U} as “worlds” in this context.

1. **Cross-World Abduction:** Update $P_B(\mathbf{U}^*, \mathbf{U})$ from (3.3) by the evidence $(\mathbf{x}^*, \mathbf{z},)$ to obtain the joint (“cross-world”) posterior $P(\mathbf{U}^*, \mathbf{U} \mid \mathbf{x}^*, \mathbf{z})$ given by

$$P_B(\mathbf{u}^*, \mathbf{u} \mid \mathbf{x}^*, \mathbf{z}) = \frac{P_B(\mathbf{u}^*, \mathbf{u})}{P_B(\mathbf{x}^*, \mathbf{z})} \mathbf{1}_{\{\mathbf{X}^*(\mathbf{u}^*)=\mathbf{x}^*\}} \mathbf{1}_{\{\mathbf{Z}(\mathbf{u})=\mathbf{z}\}} \quad (3.5)$$

where $P_B(\mathbf{x}^*, \mathbf{z})$ is given by (3.4).⁹

2. **Marginalisation:** Marginalise out \mathbf{U} to obtain the counterfactual posterior $P_B(\mathbf{U}^* \mid \mathbf{x}^*, \mathbf{z})$:

$$P_B(\mathbf{u}^* \mid \mathbf{x}^*, \mathbf{z}) = \sum_{\mathbf{u}} P_B(\mathbf{u}^*, \mathbf{u} \mid \mathbf{x}^*, \mathbf{z}). \quad (3.6)$$

3. **Prediction:** Use the model $(\mathcal{M}, P_B(\mathbf{U}^* \mid \mathbf{x}^*, \mathbf{z}))$ to predict \mathbf{Y}^* :

$$P_B(\mathbf{y}^* \mid \mathbf{x}^*, \mathbf{z}) = \sum_{\mathbf{u}^*} P_B(\mathbf{u}^* \mid \mathbf{x}^*, \mathbf{z}) \mathbf{1}_{\{\mathbf{Y}^*(\mathbf{u}^*)=\mathbf{y}^*\}} \quad (3.7)$$

Remark 3 Whereas for interventional counterfactuals the antecedent is always attainable by intervention, backtracking counterfactuals may have no solution if the antecedent is incompatible with the causal laws. For example, if $Y := X$, any backtracking counterfactual with antecedent $y^* \neq x^*$ cannot be realized. Dealing with such “counterlegals” (Fisher, 2017b) appropriately requires a semantics that combines backtracking and interventions. We return to this in Appx. B.

To illustrate this procedure, we now carry out the respective calculations for Ex. 2.

Example 2 (continued) We now compute $P_B(X^*, Z^* \mid Y^* = 3, X = 1, Y = 2, Z = 2)$. Suppose that we use the backtracking conditional $P_B(\mathbf{U}^* \mid \mathbf{U} = \mathbf{u}) = \mathcal{N}(\mathbf{u}, \text{diag}(\sigma_X^2, \sigma_Y^2, \sigma_Z^2))$ which assigns higher probability to values of \mathbf{U}^* that are close to a given \mathbf{u} , with the relative scale of closeness given by $\sigma_X^2, \sigma_Y^2, \sigma_Z^2$ (see § 3.3 for more details).

1. **Cross-World Abduction:** Ignoring multiplicative constants w.r.t. $(\mathbf{u}^*, \mathbf{u})$, we find

$$P_B(\mathbf{u}^*, \mathbf{u} \mid Y^* = 3, X = 1, Y = 2, Z = 2) \propto P_B(\mathbf{u}^*, \mathbf{u}) \mathbf{1}_{\{\mathbf{u}=(1,1,-1)\}} \mathbf{1}_{\{u_Y^*=3-u_X^*\}} \quad (3.8)$$

2. **Marginalisation:** Only one non-zero term in the sum over \mathbf{u} ’s in (3.6) remains:

$$\begin{aligned} P_B(\mathbf{u}^* \mid Y^* = 3, X = 1, Y = 2, Z = 2) &\propto P_B(\mathbf{u}^* \mid \mathbf{u} = (1, 1, -1)) \mathbf{1}_{\{u_Y^*=3-u_X^*\}} \\ &\propto \exp \left\{ -\frac{(u_X^*-1)^2}{2\sigma_X^2} - \frac{(3-u_X^*-1)^2}{2\sigma_Y^2} - \frac{(u_Z^*+1)^2}{2\sigma_Z^2} \right\} \mathbf{1}_{\{u_Y^*=3-u_X^*\}} \end{aligned} \quad (3.9)$$

Completing the square for u_X^* , we finally obtain the counterfactual posterior:

$$\mathbf{U}^* \mid (Y^*=3, X=1, Y=2, Z=2) \sim \mathcal{N}_{(U_X^*, U_Z^*)} \left(\begin{pmatrix} 1 + \frac{r}{1+r} \\ -1 \end{pmatrix}, \begin{pmatrix} \frac{\sigma_X^2}{1+r} & 0 \\ 0 & \sigma_Z^2 \end{pmatrix} \right) \delta(U_Y^*=3 - U_X^*)$$

where $r = \frac{\sigma_X^2}{\sigma_Y^2}$ denotes the ratio of variances from the backtracking conditional.

3. **Prediction:** Using the relation $(X^*, Z^*) = (U_X^*, 3 + U_X^* + U_Z^*)$ derived from the causal laws, we finally obtain the solution as a linear transformation of the Gaussian from step 2.:

$$(X^*, Z^*) \mid (Y^*=3, X=1, Y=2, Z=2) \sim \mathcal{N} \left(\begin{pmatrix} 1 + \frac{r}{1+r} \\ 3 + \frac{r}{1+r} \end{pmatrix}, \begin{pmatrix} \frac{\sigma_X^2}{1+r} & \frac{\sigma_X^2}{1+r} \\ \frac{\sigma_X^2}{1+r} & \sigma_Z^2 + \frac{\sigma_X^2}{1+r} \end{pmatrix} \right) \quad (3.10)$$

This form of (3.10) is consistent with intuition: if changes in U_X^* and U_Y^* are considered equally likely (i.e., for $\sigma_X^2 = \sigma_Y^2$ in the backtracking conditional $P_B(\mathbf{U}^* \mid \mathbf{U})$), we have $r = 1$ and case (iii) in (3.2) is the most likely (maximum a posteriori) scenario; if, on the other hand, U_X^* is considered much less likely to change (i.e., $\sigma_X^2 \ll \sigma_Y^2$ and $r \rightarrow 0$), (3.10) centers around case (i); conversely, if U_Y^* is much less likely to change (i.e., $\sigma_X^2 \gg \sigma_Y^2$ and $r \rightarrow \infty$), (3.10) centers around case (ii).

9. This step can also be viewed as the projection of $P_B(\mathbf{U}^*, \mathbf{U})$ onto the subspace consistent with $(\mathbf{x}^*, \mathbf{z})$.

3.3. On the Choice of the Backtracking Conditional $P_B(\mathbf{U}^* | \mathbf{U})$

As we have seen, $P_B(\mathbf{U}^* | \mathbf{U})$ plays an important role in answering backtracking counterfactuals: different choices may yield different answers. We therefore now discuss desiderata for, and relevant properties of, the backtracking conditional $P_B(\mathbf{U}^* | \mathbf{U})$.

First, an intuitive desideratum is that $P_B(\mathbf{U}^* | \mathbf{U})$ should assign high probability to values of \mathbf{U}^* that are considered close to a given \mathbf{u} , and small probability to those that are considered far.

Property 1 (Preference for Closeness) *We say that $P_B(\mathbf{U}^* | \mathbf{U})$ has Preference for Closeness if*

$$\forall \mathbf{u} : \arg \max_{\mathbf{u}^*} P_B(\mathbf{u}^* | \mathbf{u}) = \{\mathbf{u}\}. \quad (3.11)$$

Note that we could also impose such preference for closeness w.r.t. the mean or median, instead of (or in addition to) the mode as in (3.11). Second, to avoid a-priori asymmetries between the factual and counterfactual world, the following symmetry requirement also seems natural.

Property 2 (Symmetry) *We say that $P_B(\mathbf{U}^* | \mathbf{U})$ is symmetric if for any $(\mathbf{u}^*, \mathbf{u})$:*

$$P_B(\mathbf{u}^* | \mathbf{u}) = P_B(\mathbf{u} | \mathbf{u}^*) \quad (3.12)$$

Note that this implies and is implied by matching marginals, or priors, across worlds:

Lemma 4 (Symmetry is equivalent to matching marginals) *$P_B(\mathbf{U}^* | \mathbf{U})$ is symmetric if and only if the marginal of \mathbf{U}^* induced by $P(\mathbf{U})$ and $P_B(\mathbf{U}^* | \mathbf{U})$ matches $P(\mathbf{U})$, that is:*

$$P_B(\mathbf{U}^*) := \sum_{\mathbf{u}} P_B(\mathbf{U}^* | \mathbf{u}) P(\mathbf{u}) = P(\mathbf{U}). \quad (3.13)$$

Proof The result follows from (3.12), Bayes rule, and the equality of domains of \mathbf{U} and \mathbf{U}^* . ■

Third, since it is often assumed that the exogenous variables are mutually independent (i.e. that $P(\mathbf{U})$ factorises), a similar natural property for the backtracking conditional is that each counterfactual variable U_j^* only depends on its factual counterpart U_j .

Property 3 (Decomposability) *We say that $P_B(\mathbf{U}^* | \mathbf{U})$ is decomposable if for any $(\mathbf{u}^*, \mathbf{u})$:*

$$P_B(\mathbf{u}^* | \mathbf{u}) = \prod_{j=1}^m P_B(u_j^* | u_j). \quad (3.14)$$

In line with the philosophical notion of *closest possible worlds*, one option is to construct $P_B(\mathbf{U}^* | \mathbf{U})$ based on a distance function $d(\cdot, \cdot)$ defined over the shared domain of \mathbf{U} and \mathbf{U}^* :

$$P_B(\mathbf{u}^* | \mathbf{u}) = \frac{1}{Z} \exp\{-d(\mathbf{u}^*, \mathbf{u})\} \quad (3.15)$$

where $Z = \sum_{\mathbf{u}^*} \exp\{-d(\mathbf{u}^*, \mathbf{u})\}$ is a normalisation constant.

Lemma 5 *Any $P_B(\mathbf{u}^* | \mathbf{u})$ of the form (3.15) satisfies Properties 1 and 2. Further, if $d(\mathbf{u}^*, \mathbf{u})$ can be written as $\sum_j d(u_j^*, u_j)$, then $P_B(\mathbf{u}^* | \mathbf{u})$ also satisfies Property 3.*

Proof By definition, any distance satisfies $d(\mathbf{u}^*, \mathbf{u}) \geq 0$ with equality iff $\mathbf{u}^* = \mathbf{u}$ (implying Property 1) and $d(\mathbf{u}^*, \mathbf{u}) = d(\mathbf{u}, \mathbf{u}^*)$ (implying Property 2). Property 3 follows from substituting the sum of component-wise distances into (3.15) and writing it as a product of exponentials. ■

For real valued $\mathbf{U} \in \mathbb{R}^m$, a natural choice is the (squared) Mahalanobis distance

$$d(\mathbf{u}^*, \mathbf{u}) = \frac{1}{2}(\mathbf{u} - \mathbf{u}^*)^\top \Sigma^{-1}(\mathbf{u} - \mathbf{u}^*) \quad (3.16)$$

for some positive-definite, symmetric (covariance) matrix Σ , giving rise to the multivariate *Gaussian backtracking conditional* $P_B(\mathbf{U}^* | \mathbf{U} = \mathbf{u}) = \mathcal{N}(\mathbf{u}, \Sigma)$ (for which mode, median and mean coincide). If $\Sigma = \text{diag}(\sigma_1^2, \dots, \sigma_m^2)$ as used in Ex. 2, this satisfies Properties 1, 2, and 3.

We emphasize that the distance-based (3.15) is but one option for specifying P_B . A simple alternative is to dismiss entirely what actually happened and take $P_B(\mathbf{u}^* | \mathbf{u}) = P(\mathbf{u}^*)$, corresponding to the extreme view that what would have happened is completely independent of the factual events, and is instead determined entirely by the prior probability. More generally, one could combine both options by using parameters $\alpha, \beta \geq 0$ to weigh off the relative importance of prior and distance,

$$P_B(\mathbf{u}^* | \mathbf{u}) = \frac{1}{Z} P(\mathbf{u}^*)^\alpha \exp\{-\beta d(\mathbf{u}^*, \mathbf{u})\}, \quad (3.17)$$

where $Z = \sum_{\mathbf{u}^*} P(\mathbf{u}^*)^\alpha \exp\{-\beta d(\mathbf{u}^*, \mathbf{u})\}$ is a normalisation constant.

3.4. Theoretical Insights

Our first theoretical insight is that exogenous non-ancestors of factual and counterfactual observations remain unaffected, in that their posterior is equal to their prior.

Proposition 6 *Let \mathcal{M} be an acyclic causal model. Suppose that $P(\mathbf{U})$ factorises and that $P_B(\mathbf{U}^* | \mathbf{U})$ is decomposable. For any $\mathbf{X}, \mathbf{Z} \subseteq \mathbf{V}$ with (endogenous and exogenous) ancestors $\text{AN}_{\mathbf{X}}, \text{AN}_{\mathbf{Z}}$:*

$$\mathbf{U}_J \subseteq \mathbf{U} \setminus (\text{AN}_{\mathbf{X}} \cup \text{AN}_{\mathbf{Z}}) \implies \forall (\mathbf{x}^*, \mathbf{z}) : P(\mathbf{U}_J^*, \mathbf{U}_J | \mathbf{x}^*, \mathbf{z}) = \prod_{j \in J} P(U_j^*, U_j) \quad (3.18)$$

Proof The result follows from d-separation between $(\mathbf{U}_J^*, \mathbf{U}_J)$ and $(\mathbf{X}^*, \mathbf{Z})$. \blacksquare

Our second insight concerns the distinguishability of SCMs based on backtracking counterfactuals. As is well-known, identifying probabilities of *interventional* counterfactuals requires full knowledge of the laws \mathbf{F} , which is hard to obtain in practice (Bareinboim et al., 2022). Interestingly, probabilities of *observational* counterfactuals are not as demanding since they only depend on the solution or “reduced form” $\mathbf{V}(\mathbf{U})$ (see, e.g., Schölkopf and von Kügelgen, 2022, § 10 for details).

Proposition 7 *Let $(\mathcal{M}_1, P(\mathbf{U}))$ and $(\mathcal{M}_2, P(\mathbf{U}))$ be causal models over the same variables \mathbf{V} and \mathbf{U} whose laws \mathbf{F}_1 and \mathbf{F}_2 have identical solutions $\mathbf{V}_1(\mathbf{u}) = \mathbf{V}_2(\mathbf{u})$ for all values \mathbf{u} . Then for any choice of $P_B(\mathbf{U}^* | \mathbf{U})$, both models will imply identical $P_B(\mathbf{y}^*, \mathbf{z})$ for all choices \mathbf{y}^* and \mathbf{z} .*

Proof This is an immediate consequence of (3.4) only involving (3.3) and the solution function. \blacksquare

In particular, this result implies that causal structure is not discernible purely based on backtracking counterfactuals, as demonstrated by the following example.

Example 3 *Consider the following three causal models over $\mathbf{V} = \{X, Y\}$ and $\mathbf{U} = \{U\}$ with laws $\mathbf{F}_{X \rightarrow Y} = \{X := U, Y := X\}$, $\mathbf{F}_{Y \rightarrow X} = \{Y := U, X := Y\}$, $\mathbf{F}_{X \leftrightarrow Y} = \{X := U, Y := U\}$, and some shared $P(\mathbf{U})$. Then all share the same reduced form $(X, Y) = (U, U)$ and hence the same backtracking counterfactuals, despite differing in their causal diagrams.*

4. Related Work

Most closely related to our formalisation is the Extended Structural Model proposed by Lucas and Kemp (2015, Eq. 1) to model recent empirical findings on the context-dependency of humans’ use of backtracking. It employs a decomposable backtracking conditional for Boolean variables,

$$P_B(U_j^* | U_j) = s\delta(U_j) + (1 - s)P(U_j^*), \quad (4.1)$$

where the “stability” hyperparameter $s \in [0, 1]$ interpolates between independent worlds ($s = 0$) and perfectly shared exogenous variables ($s = 1$). However, forcing \mathbf{U}^* to either copy \mathbf{U} or to completely ignore it is quite restrictive, see (3.17) for a more flexible implementation of this idea.

Other existing accounts are less formal: they mostly consider non-probabilistic Boolean conditionals of the form “Given world w , if \mathbf{x}^* were true then \mathbf{y}^* would be true” and typically involve minimising the number of (endogenous and) exogenous variables that change across worlds. Hiddleston (2005)’s semantics focuses on minimising the number of exogenous non-descendants of the antecedent that change values, which corresponds to using a variant of the Manhattan distance in (3.15). Interestingly, Pearl’s interventional semantics (§ 2) can be interpreted in a similar vein, instead minimising the number of intervened-upon variables. In fact, he uses this formulation to argue that this is essentially equivalent to Lewis’s possible world semantics (Pearl, 2009, p.241). Fisher (2017b) combines both approaches by applying Hiddleston’s condition if there exists a backtracking solution and resorting to a minimal number of interventions only if necessary. Lee (2017) also considers both semantics, but does not combine them into a single one.

Empirical research by cognitive scientists confirms that depending on the context, humans interpret counterfactuals as either backtracking or non-backtracking: “participants are more likely to backtrack when explicitly asked to consider a counterfactual’s *causes*. However, when directly asked about the *effects* of a counterfactual state, most people don’t backtrack” (Gerstenberg et al., 2013). Moreover, people use the broader context, exact wording, and level of determinism of the involved mechanism “to infer how the antecedent is most likely to have come about” and use this to decide whether to backtrack or not (Rips and Edwards, 2013). It has also been suggested that backtracking is preferred if and only if doing so makes a counterfactual claim true (Han et al., 2014).

When evaluating counterfactual explanations for specific historical events, historians rely on the *minimal rewrite rule*. This rule requires making minimal changes to the actual world to create the necessary conditions that would have led to a particular counterfactual antecedent (Tetlock and Belkin, 1996). Reiss has argued that historians interpret minimality in a backtracking sense, stating that “the antecedent is not implemented by a miracle”, but rather “counterfactuals in history are backtracking” (Reiss, 2009, p. 719). His informal description closely aligns with our formal semantics, as he states that counterfactuals are evaluated based on “causal generalizations” that must not be violated, and only those counterfactual background conditions that “were likely” are to be considered. Additionally, using models similar to those from Ex. 3, Reiss observes that these historical counterfactuals cannot distinguish between the causal structure of a chain of causes and a common cause structure, a finding that is entirely consistent with our Proposition 7.

5. Connections and Applications to Explainable AI (XAI)

We believe that backtracking counterfactuals may hold great promise for XAI, which is concerned with offering explanations for the decisions of a machine learning model implementing a function $Y = f(\mathbf{X})$, where \mathbf{X} are *input features* and Y is the *output or label*. Given that we observe some input \mathbf{x} and corresponding output y , the general aim is to find a feature subset $\mathbf{Z} \subseteq \mathbf{X}$ that “explains” $y = f(\mathbf{x})$. A particularly promising approach has been to look for so-called nearest *counterfactual explanations*, meaning that we look for both \mathbf{z} and \mathbf{z}^* such that changing \mathbf{Z} from \mathbf{z} to \mathbf{z}^* would have resulted in $y^* \neq y$, and \mathbf{z} and \mathbf{z}^* are close according to some distance function (Wachter et al., 2017). Although the general idea is rather intuitive, there has been much discussion on how exactly to implement it. The core problem lies in the choice of values \mathbf{w}^* for the remaining features $\mathbf{W}^* = \mathbf{X}^* \setminus \mathbf{Z}^*$ that should accompany \mathbf{z}^* so that $f(\mathbf{z}^*, \mathbf{w}^*) = y^*$.

This discussion between AI-researchers can in fact be seen as the operational analogue of the philosophical discussion on the semantics of counterfactuals that we started out with in § 1, and which is far from resolved. Hence, we are faced with the same dilemma as before: to backtrack, or not to backtrack? Interestingly, the proposal of Wachter et al. (2017) does neither, keeping all other variables fixed at their factual values $\mathbf{w}^* = \mathbf{w}$. This proposal has recently come under criticism for its failure to take into account causal dependencies. On the one hand, Beckers (2022) criticises it for failing to follow Pearl (and the entire causal modelling tradition) in maintaining an *interventionist* reading of counterfactuals. On the other, Mahajan et al. (2019); Crupi et al. (2022) criticise it for ignoring violations of the causal laws, and offer alternative proposals that—without making it explicit—can be viewed as building on a kind of *backtracking* semantics instead, though ones that involve distances between endogenous variables or require knowledge of the causal laws.

When counterfactual explanations are offered in the context of *algorithmic recourse* (Ustun et al., 2019), meaning that one is looking for actionable changes to the input features that result in a more favourable outcome, the first criticism is undoubtedly on the right track (Karimi et al., 2020, 2021). For example, if an agent is told to increase their income by \$5,000 in order to be granted a loan, then this explanation only succeeds in correctly predicting the decision if it has correctly taken into account the forward-tracking consequences of this change on any features that are downstream of income, such as savings. If these changes are implemented externally (thereby overriding the observational distribution), they should indeed be modelled as interventions.

Yet, in most cases the use of interventional counterfactuals is unrealistic, since they cannot be identified solely from the observational distribution that ML-methods are trained on. Moreover, it is not at all clear that interventional counterfactuals are the right approach in the context of explanations that aim to *contest*, or to simply *understand* or *diagnose*, the outcome that was reached. In line with the second criticism (and the recent trend in the philosophy literature), we claim that backtracking counterfactuals deserve attention as a promising alternative for these contexts, and we have here developed the tools to do so. Importantly, as Proposition 7 suggests, backtracking counterfactuals should be easier to identify, as they depend on the causal dependencies between endogenous variables only through the reduced form. By leaving the causal laws intact, they also remain closer to the observational distribution, or data manifold, which is a desirable constraint for counterfactual explanations (Wexler et al., 2019; Joshi et al., 2019; Poyiadzi et al., 2020; Sharma et al., 2020).

An implementation of this idea could look roughly as follows. Let $(\mathcal{M}, P(\mathbf{U}))$ be a probabilistic causal model with endogenous variables $\mathbf{V} = \mathbf{X} \cup \{Y\}$, $P(\mathbf{U})$ having full support, and laws such that $Y = f(\mathbf{X})$ holds with probability one. For a particular choice of $P_B(\mathbf{U}^* | \mathbf{U})$, we then say that “ \mathbf{x} rather than \mathbf{x}^* explains why $f(\mathbf{x}) = y$ rather than $y^* \neq y$ ” if such a change to y^* would be most likely to have come about through \mathbf{x}^* , that is, if $\mathbf{x}^* \in \arg \max_{\mathbf{x}^*} P_B(\mathbf{x}^* | y^*, \mathbf{x}, y)$. Note that this is exactly the *diagnostic* kind of reasoning that only backtracking allows, thus establishing a clear link between counterfactual explanations and maximum a posteriori backtracking counterfactuals.

In practice, interpretable explanations should be concise, highlighting just a few features that are very likely to have been different. To implement this, we could set a probability threshold α and a maximal number k of “explanatory” features, and then look for an optimal feature subset $\mathbf{Z} \subseteq \mathbf{X}$ with $|\mathbf{Z}| \leq k$ and altered counterfactual feature values \mathbf{z}^* satisfying $z_i^* \neq z_i$ for all $Z_i \in \mathbf{Z}$, such that $P_B(\mathbf{z}^* | y^*, \mathbf{x}, y) > \alpha$.¹⁰ Lastly, note that a variant of the original proposal can be retrieved within our proposed backtracking framework by considering $\arg \max_{\mathbf{z}^*} P_B(\mathbf{z}^* | \mathbf{W}^* = \mathbf{w}, y^*, \mathbf{x}, y)$ for all choices $\mathbf{Z} \subseteq \mathbf{X}$, where as before $\mathbf{X} = \mathbf{W} \cup \mathbf{Z}$ and $z_i^* \neq z_i$ for all $Z_i \in \mathbf{Z}$.

10. For continuous variables, we could, for example, instead set a distance threshold ϵ and then look for a set A of values that are at least distance ϵ away, $d(\mathbf{z}^*, \mathbf{z}) > \epsilon$, and satisfy $P_B(\mathbf{z}^* \in A | y^*, \mathbf{x}, y) > \alpha$.

Causal Attribution Analysis. Counterfactual explanations are also particularly relevant for causal attribution tasks such as *root cause analysis* of an outlier event $Y = y$ (Budhathoki et al., 2022). A central idea is that the exogenous variables in an SCM (the “roots”) ultimately explain why $Y = y$. The proposed method thus involves *keeping the causal laws intact* while varying the values of some exogenous variables according to a “counterfactual distribution”, which is emphasised as “the key ingredient” (Budhathoki et al., 2022, p. 4). This can be viewed as a form of backtracking, albeit not formulated in those terms. Suppose we observe an outlier $Y = y$ together with $\mathbf{X} = \mathbf{x}$, and—assuming invertibility of the reduced form $\mathbf{V}(\mathbf{U})$ —use this to infer $\mathbf{U} = \mathbf{u}$. The main building block of Budhathoki et al.’s method is to quantify how much each $\mathbf{U}_i \subseteq \mathbf{U}$ contributed to the outlier event by computing the counterfactual probability of a similar (or more extreme) outlier, conditional on the factual $\mathbf{U} = \mathbf{u}$ and on a subset $\mathbf{U}_S \subset \mathbf{U}$ being fixed across worlds while the remaining $\mathbf{U}_{-S} = \mathbf{U} \setminus \mathbf{U}_S$ are resampled from the prior $P(\mathbf{U})$, combined with Shapley-value based symmetrization. We can give a precise expression to this as $P_B(\tau(y^*) \geq \tau(y) \mid \mathbf{U}_S^* = \mathbf{u}_S, \mathbf{U} = \mathbf{u})$ with backtracking conditional $P_B(\mathbf{U}^* \mid \mathbf{U}) = P(\mathbf{U}^*)$, where τ is a calibration function to allow for comparing outliers over different ranges. Similar ideas based on “structure-preserving interventions” that do not change the laws \mathbf{F} but only resample subsets of the exogenous \mathbf{U} have also been used by Janzing et al. (2021) for attributing quantities such as entropy or variance to other variables in the model.

6. Conclusion and Future Work

We have here presented the first formal account of backtracking counterfactuals within probabilistic causal models (§ 3). Doing so in full generality required the introduction of a new object, the backtracking conditional, which quantifies a notion of similarity between worlds (§ 3.2). This involves a design choice that can take into account the specific context and the modeller’s background assumptions. In fact, many previous accounts of backtracking can be seen as specific choices of such a conditional (§ 4). We have laid out some sensible desiderata and means to achieve them in § 3.3.

As stressed throughout, we do not view backtracking as a replacement of the interventionist account (§ 2) but rather as complementary to it. Our work thus also emphasizes the ambiguous nature of counterfactuals and their semantics, a point which we discuss further through a worked-out example in Appx. C. Proposition 7 and Ex. 3 make clear that backtracking counterfactuals do not allow for discerning causal structure and are therefore of limited use for interventional reasoning. At the same time, they are helpful for the kind of diagnostic reasoning that occurs in settings where interventions are inconsistent with human judgement (§ 4), require knowledge that is unavailable (§ 5), or are perhaps not even meaningful. The latter shows up in causal fairness analysis (Kusner et al., 2017; Kilbertus et al., 2017; von Kügelgen et al., 2022; Plecko and Bareinboim, 2022), which often focuses on such attributes (like race or gender), leading some to dispute their status as valid causes (Holland, 2008; Hu and Kohler-Hausmann, 2020). For this reason, we consider applications of backtracking for fairness analysis an interesting future direction.

By endowing backtracking with a formal semantics within the same general SCM framework also used for interventional counterfactuals, the present work paves the way for a unified framework of both backtracking *and* non-backtracking counterfactual reasoning. We make a proposal for incorporating hard interventions in Appx. B, but future work is needed to provide a more comprehensive and rigorous account. The general idea is to allow modifying both the causal laws and the background conditions, and to weigh off changes to both of them through an appropriate choice of backtracking conditional. This would not only imply the property of being guaranteed a solution (recall Remark 3), but may also allow for more accurate models of human counterfactual reasoning.

Acknowledgments

The authors thank Dominik Janzing for insightful discussions, and the anonymous reviewers and the area chair for helpful comments and suggestions. This work was supported by the Tübingen AI Center and the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany’s Excellence Strategy – EXC number 2064/1 – Project number 390727645.

References

- Alexander Balke and Judea Pearl. Probabilistic evaluation of counterfactual queries. In *Proceedings of the 12th National Conference on Artificial Intelligence*, pages 230–237. AAAI Press / The MIT Press, 1994. [Cited on pages [4](#) and [5](#).]
- 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*, pages 507–556, 2022. [Cited on pages [10](#) and [17](#).]
- Sander Beckers. Causal explanations and XAI. In *Proceedings of the First Conference on Causal Learning and Reasoning*, volume 177, pages 90–109. PMLR, 2022. [Cited on page [12](#).]
- Kailash Budhathoki, Lenon Minorics, Patrick Blöbaum, and Dominik Janzing. Causal structure-based root cause analysis of outliers. In *International Conference on Machine Learning*, pages 2357–2369. PMLR, 2022. [Cited on page [13](#).]
- Riccardo Crupi, Alessandro Castelnovo, Daniele Regoli, and Beatriz San Miguel Gonzalez. Counterfactual explanations as interventions in latent space. *Data Mining and Knowledge Discovery*, pages 1–37, 2022. [Cited on page [12](#).]
- Philip Dawid. Causal inference without counterfactuals. *Journal of the American statistical Association*, 95(450):407–424, 2000. [Cited on page [4](#).]
- Cian Dorr. Against counterfactual miracles. *The Philosophical Review*, 125(2):241–286, 2016. [Cited on page [2](#).]
- Michael Esfeld. Super-humeanism and free will. *Synthese*, 198(7):6245–6258, 2021. [Cited on page [2](#).]
- Tyrus Fisher. Causal counterfactuals are not interventionist counterfactuals. *Synthese*, 194(12):4935–4957, 2017a. [Cited on page [2](#).]
- Tyrus Fisher. Counterlegal dependence and causation’s arrows: Causal models for backtrackers and counterlegals. *Synthese*, 194(12):4983–5003, 2017b. [Cited on pages [2](#), [8](#), and [11](#).]
- Tobias Gerstenberg, Christos Bechlivanidis, and David A Lagnado. Back on track: Backtracking in counterfactual reasoning. In *Proceedings of the Annual Meeting of the Cognitive Science Society*, volume 35, 2013. [Cited on pages [2](#) and [11](#).]
- Jung-Ho Han, William Jimenez-Leal, and Steve Sloman. Conditions for backtracking with counterfactual conditionals. In *Proceedings of the Annual Meeting of the Cognitive Science Society*, volume 36, 2014. [Cited on page [11](#).]

- Eric Hiddleston. A causal theory of counterfactuals. *Noûs*, 39(4):632–657, 2005. [Cited on page 11.]
- Paul W Holland. Causation and race. *White logic, white methods: Racism and methodology*, pages 93–109, 2008. [Cited on page 13.]
- Lily Hu and Issa Kohler-Hausmann. What’s sex got to do with machine learning? In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, pages 513–513, 2020. [Cited on page 13.]
- Dominik Janzing, Patrick Blöbaum, Lenon Minorics, Philipp Faller, and Atalanti Mastakouri. Quantifying intrinsic causal contributions via structure preserving interventions. *arXiv 2007.00714*, 2021. [Cited on page 13.]
- Shalmali Joshi, Oluwasanmi Koyejo, Warut Vijitbenjaronk, Been Kim, and Joydeep Ghosh. Towards realistic individual recourse and actionable explanations in black-box decision making systems. *arXiv preprint arXiv:1907.09615*, 2019. [Cited on page 12.]
- Amir-Hossein Karimi, Julius Von Kügelgen, Bernhard Schölkopf, and Isabel Valera. Algorithmic recourse under imperfect causal knowledge: a probabilistic approach. *Advances in Neural Information Processing Systems*, 33:265–277, 2020. [Cited on page 12.]
- 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*, pages 353–362, 2021. [Cited on page 12.]
- Niki Kilbertus, Mateo Rojas Carulla, Giambattista Parascandolo, Moritz Hardt, Dominik Janzing, and Bernhard Schölkopf. Avoiding discrimination through causal reasoning. *Advances in neural information processing systems*, 30, 2017. [Cited on page 13.]
- Matt J Kusner, Joshua Loftus, Chris Russell, and Ricardo Silva. Counterfactual fairness. *Advances in neural information processing systems*, 30, 2017. [Cited on page 13.]
- Kok Yong Lee. Hiddleston’s causal modeling semantics and the distinction between forward-tracking and backtracking counterfactuals. *Studies in Logic*, 10(1), 2017. [Cited on page 11.]
- David Lewis. *Counterfactuals*. Oxford: Blackwell Publishers and Cambridge, MA: Harvard University Press, 1973. [Cited on page 1.]
- David Lewis. Counterfactual dependence and time’s arrow. *Noûs*, pages 455–476, 1979. [Cited on pages 1 and 11.]
- Barry Loewer. Counterfactuals and the second law. In Huw Price and Richard Corry, editors, *Causation, Physics, and the Constitution of Reality: Russell’s Republic Revisited*. Oxford University Press, 2007. [Cited on page 2.]
- Barry Loewer. The consequence argument meets the mentaculus, 2020. Working papers, Rutgers University. [Cited on page 2.]
- Christopher G Lucas and Charles Kemp. An improved probabilistic account of counterfactual reasoning. *Psychological review*, 122(4):700, 2015. [Cited on pages 2 and 10.]

- Divyat Mahajan, Chenhao Tan, and Amit Sharma. Preserving causal constraints in counterfactual explanations for machine learning classifiers. *arXiv:1912.03277*, 2019. [Cited on page 12.]
- Judea Pearl. *Causality*. Cambridge university press, 2009. [Cited on pages 1, 2, 3, 4, 11, and 18.]
- Judea Pearl and Dana Mackenzie. *The book of why: the new science of cause and effect*. Basic books, 2018. [Cited on page 17.]
- Drago Plecko and Elias Bareinboim. Causal fairness analysis. *arXiv preprint arXiv:2207.11385*, 2022. [Cited on page 13.]
- Rafael Poyiadzi, Kacper Sokol, Raul Santos-Rodriguez, Tijl De Bie, and Peter Flach. Face: feasible and actionable counterfactual explanations. In *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*, pages 344–350, 2020. [Cited on page 12.]
- Julian Reiss. Counterfactuals, thought experiments, and singular causal analysis in history. *Philosophy of Science*, 76(5):712–723, 2009. [Cited on pages 2 and 11.]
- Lance J Rips. Two causal theories of counterfactual conditionals. *Cognitive science*, 34(2):175–221, 2010. [Cited on page 2.]
- Lance J Rips and Brian J Edwards. Inference and explanation in counterfactual reasoning. *Cognitive Science*, 37(6):1107–1135, 2013. [Cited on page 11.]
- Bernhard Schölkopf and Julius von Kügelgen. From statistical to causal learning. *arXiv preprint arXiv:2204.00607*, 2022. [Cited on page 10.]
- Shubham Sharma, Jette Henderson, and Joydeep Ghosh. Certifai: A common framework to provide explanations and analyse the fairness and robustness of black-box models. In *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*, pages 166–172, 2020. [Cited on page 12.]
- Philip E Tetlock and Aaron Belkin. *Counterfactual thought experiments in world politics: Logical, methodological, and psychological perspectives*. Princeton University Press, 1996. [Cited on page 11.]
- Berk Ustun, Alexander Spangher, and Yang Liu. Actionable recourse in linear classification. In *Proceedings of the conference on fairness, accountability, and transparency*, pages 10–19, 2019. [Cited on page 12.]
- Julius von Kügelgen, Amir-Hossein Karimi, Umang Bhatt, Isabel Valera, Adrian Weller, and Bernhard Schölkopf. On the fairness of causal algorithmic recourse. In *Proceedings of the 36th AAAI Conference on Artificial Intelligence*, 2022. [Cited on page 13.]
- Sandra Wachter, Brent Mittelstadt, and Chris Russell. Counterfactual explanations without opening the black box: Automated decisions and the GDPR. *Harv. JL & Tech.*, 31:841, 2017. [Cited on pages 2, 11, and 12.]
- James Wexler, Mahima Pushkarna, Tolga Bolukbasi, Martin Wattenberg, Fernanda Viégas, and Jimbo Wilson. The what-if tool: Interactive probing of machine learning models. *IEEE transactions on visualization and computer graphics*, 26(1):56–65, 2019. [Cited on page 12.]
- James Woodward. *Causation with a human face: Normative theory and descriptive psychology*. Oxford University Press, 2021. [Cited on pages 2 and 3.]

Appendix A. Ladder of Causation

The interventional counterfactual distribution $P((X^*, Z^*)_{Y^*=3} \mid X = 1, Y = 2, Z = 2)$ of interest in Ex. 1 differs from both simple conditioning, that is, the observational distribution

$$P(X, Z \mid Y = 3) = \mathcal{N} \left(\begin{pmatrix} 1.5 \\ 4.5 \end{pmatrix}, \begin{pmatrix} 0.5 & 0.5 \\ 0.5 & 0.5 \end{pmatrix} \right), \quad (\text{A.1})$$

as well as from the interventional distribution

$$P((X, Z)_{Y=3}) = \mathcal{N} \left(\begin{pmatrix} 0 \\ 3 \end{pmatrix}, \begin{pmatrix} 1 & 1 \\ 1 & 2 \end{pmatrix} \right), \quad (\text{A.2})$$

both of which preserve uncertainty about (X, Z) . In contrast, the interventional counterfactual is fully determined in Ex. 1, because the shared background conditions \mathbf{u} could be uniquely inferred from the factual observation. These three modes of reasoning—observational, interventional, and counterfactual—constitute increasingly difficult tasks, each requiring additional data or assumptions over the previous, and form the three rungs of the so-called “Ladder of Causation” (Pearl and Mackenzie, 2018) or “Pearl Causal Hierarchy” (Bareinboim et al., 2022).

Appendix B. Toward A Unified Framework of Counterfactual Reasoning

As alluded to in § 6, combining backtracking and interventional counterfactuals into a single unified framework of counterfactual reasoning is an interesting problem, and we believe that our framework is suitable for making progress toward this goal. Below, we present a first attempt at doing so. Specifically, we show how hard interventions can be modelled through auxiliary variables, a change in the causal laws, and an appropriate choice of backtracking conditional.¹¹

Let $(\mathcal{M}, P(\mathbf{U}))$ be a probabilistic causal model with laws

$$V_i := f_i(\mathbf{PA}_i, \mathbf{U}_i) \quad i = 1, \dots, n. \quad (\text{B.1})$$

Suppose that V_i takes values in $\mathcal{V}_i \subseteq \mathbb{R}$ and that we wish to reason about hard interventions of the form $\text{do}(V_i = \tilde{v}_i)$ for some constant $\tilde{v}_i \in \mathcal{V}_i$.

First, we introduce an auxiliary regime variable R_i taking values in $\mathcal{V}_i \cup \{\text{obs}\}$, where obs is a special place holder symbol that we use to denote the observational regime.

Next, we replace the i^{th} law in (B.1) with the following augmented law:

$$V_i := \tilde{f}_i(\mathbf{PA}_i, \mathbf{U}_i, R_i) = \begin{cases} f_i(\mathbf{PA}_i, \mathbf{U}_i) & \text{if } R_i = \text{obs} \\ R_i & \text{otherwise,} \end{cases} \quad (\text{B.2})$$

that is, V_i is determined by its original law in the observational regime ($R_i = \text{obs}$), and equal to R_i outside it. Note that the latter precisely captures hard interventions through $R_i = \tilde{v}_i$.

Now suppose that we wish to reason about a counterfactual involving $V_i^* = v_i^*$ and that we want to encode the notion that this counterfactual value came about through an intervention. With the above construction, this simply amounts to conditioning on the event $R_i^* = v_i^*$, and then backtracking without further modifications as explained in the main text.

More generally, to combine both backtracking and interventional semantics, we can proceed as follows: First, we introduce an auxiliary R_i and modify the causal laws as in (B.2) for each i . Then,

11. We thank the area chair for suggesting this approach.

we consider the augmented set of exogenous variables $\mathbf{U} \cup \mathbf{R}$ with $\mathbf{R} = \{R_1, \dots, R_n\}$, and place a prior on \mathbf{R} . For example, if we are interested in queries for which the factual world is purely observational, we can choose a $P(\mathbf{R})$ that (independently of \mathbf{U}) puts all mass on $\{\text{obs}\}$ for all i (though other choices are, of course, also possible). Finally, we specify a backtracking conditional $P_B(\mathbf{U}^*, \mathbf{R}^* \mid \mathbf{U}, \mathbf{R})$ in which we may encode any available background knowledge on which of the causal laws are more likely to be violated through an intervention and which are more likely to remain intact. We can then evaluate expressions of the form

$$P(\mathbf{Y}^* = \mathbf{y}, \mathbf{Z} = \mathbf{z} \mid \mathbf{R}_j^* = \mathbf{r}_j^*, \mathbf{R}_k = \mathbf{r}_k)$$

where $\mathbf{R}_j^*, \mathbf{R}_k \subseteq \mathbf{R}$ (possibly empty) specify the assumed observational or interventional regimes for some of the factual and counterfactual variables.

Note that a standard interventional counterfactual such as $P(Y_x = y \mid Z = z)$ can be retrieved as

$$P(Y^* = y \mid Z = z, R_1^* = \text{obs}, \dots, R_X^* = x, \dots, R_n^* = \text{obs}, R_1 = \text{obs}, \dots, R_n = \text{obs}),$$

and similarly for more complex conditions. That is, we specify values for all the auxiliary variables by setting them to `obs` for all factual variables and all counterfactual variables that are not intervened upon, and setting them to the appropriate intervention values for the other, intervened-upon variables (here only X). When combined with choosing $P_B(\mathbf{u}^* \mid \mathbf{u}) = \delta(\mathbf{u})$ so that the values of \mathbf{U} are forced to be identical across worlds, we recover interventional counterfactuals as a special case in this unified framework.

At the same time, we can of course also retrieve backtracking counterfactuals by setting all auxiliary variables, both counterfactual and factual, to `obs` and choosing whatever P_B one likes. Therefore, our approach offers a unified framework for interpreting counterfactuals in SCMs.

Appendix C. Firing Squad Example: Backtracking vs Interventions

We now illustrate the difference between interventional and backtracking approaches to interpreting and computing counterfactuals for the well-known firing squad example (Pearl, 2009, § 7.1.2).

C.1. Setting

Unlike the examples considered in the main paper, the firing squad is a case of a deterministic causal model with binary variables. The setting is as follows: A captain C of two riflemen A and B is waiting for a court order U on whether a prisoner P should be executed. If the court orders the execution ($U = 1$), the captain signals ($C = 1$), the two riflemen shoot ($A = B = 1$), and the prisoner dies ($P = 1$). More formally, we can express this scenario through the following SCM:

$$C := U \tag{C.1}$$

$$A := C \tag{C.2}$$

$$B := C \tag{C.3}$$

$$P := A \vee B \tag{C.4}$$

where $U \sim \text{Bernoulli}(\theta)$ without loss of generality.

Suppose that we observe that $C = A = B = P = 1$, that is, the captain signalled, both riflemen shot, and the prisoner died. We now wish to answer the following counterfactual question:

What would have happened, had rifleman A not shot?

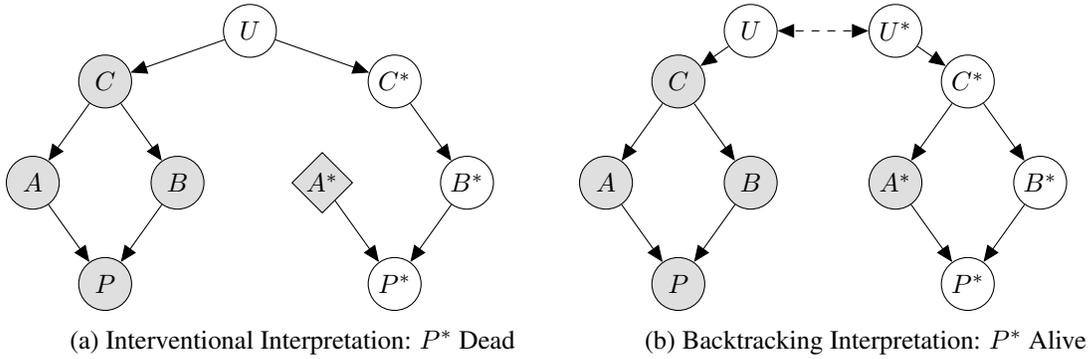


Figure 3: The Firing Squad Example Highlights the Ambiguity of Counterfactuals. We observe that the captain signals ($C = 1$), the two riflemen shoot ($A = B = 1$), and the prisoner dies ($P = 1$). We want to answer the counterfactual question: *what would have happened, had rifleman A not shot ($A^* = 0$)?* (a) The interventional interpretation is that the court still would have ordered the execution ($U = 1$) and the captain still would have signalled ($C^* = 1$) in the counterfactual world, but rifleman A did not shoot ($A^* = 0$) because they disobeyed the order or their rifle got jammed, thus locally violating the causal laws as indicated by the missing edge $C^* \rightarrow A^*$. Yet, rifleman B still would have shot ($B^* = 1$), and so the prisoner would still be dead ($P^* = 0$). (b) In stark contrast, the backtracking interpretation is that, since the causal laws must remain untouched, $A^* = 0$ could have only happened in the counterfactual world if the captain had not signalled ($C^* = 0$) because the court did not order the execution ($U^* = 0$). Hence, rifleman B also would not have shot ($B^* = 0$) and so the prisoner would still be alive ($P^* = 1$).

C.2. Interventional Counterfactual.

The interventional interpretation of our counterfactual question is illustrated in Fig. 3a. Here, we use the factual observation $C = 1$ together with (C.1) for abduction and conclude that $U = 1$. Since the exogenous variables are assumed shared in the interventionist account, we conclude that the court still would have ordered the execution in the counterfactual world. We then modify (C.2) to enforce the counterfactual antecedent $A^* = 0$ through an intervention that locally violates the causal laws. This is indicated by the missing edge $C^* \rightarrow A^*$ in Fig. 3a. A possible interpretation is that rifleman A did not shoot because they disobeyed the order of the captain or their rifle got jammed. Finally, we use the resulting modified SCM for prediction to conclude that $C^* = B^* = P^* = 1$, meaning that the captain still would have signalled, rifleman B still would have shot, and so the prisoner would still be dead.

C.3. Backtracking Counterfactual

The backtracking interpretation of our counterfactual question is illustrated in Fig. 3b. Recall that in backtracking, we introduce new counterfactual exogenous variables U^* which absorb differences across worlds, while the causal laws remain unchanged. Here, there are exactly two configurations of the variables that are compatible with the causal laws: $U = C = A = B = P = 1$ and $U = C = A = B = P = 0$, that is, all variables are either one or zero. From observing the counterfact $A^* = 0$, we conclude that we must be in the latter case. In other words, rifleman A not firing could have only happened in the counterfactual world if the captain had not signalled

($C^* = 0$) because the court did not order the execution ($U^* = 0$). Hence, rifleman B also would not have shot ($B^* = 0$) and so the prisoner would still be alive ($P^* = 0$).

C.4. Take-Away

The point of the previous example is not to single out one interpretation as correct and the other as incorrect. Rather, our goal is to highlight the ambiguous nature of counterfactuals in an intuitive context: depending on the used semantics, the same counterfactual question can be answered in radically different ways (the prisoner is either dead or alive). As discussed in § 4, depending on the circumstances, background knowledge, and prior beliefs, counterfactuals can be interpreted differently. For example, a historian following the “minimal rewrite rule” might use their background knowledge and domain understanding to answer the question interventionally if rifleman A not shooting because of disobeying the command constitutes a smaller perturbation to history than the alternative (backtracking) explanation that they perfectly obey orders and did not shoot because the captain did not signal. In the latter case, a similar decision would then need to be made as to whether the captain disobeyed the court’s order (interventional) or whether the court did not order the execution (backtracking). How a given counterfactual query is interpreted thus often depends on further background knowledge that is not made explicit in the causal model or counterfactual query. Interventional counterfactuals can be viewed as relying on an extreme form of such knowledge by always opting for a local violation of the laws (a small miracle), while leaving it implicit that they are, in fact, committing to this extreme. We offer a formal representation that allows for making this knowledge explicit and for considering different choices, such as full or partial backtracking as alternatives.