
Counterfactual Off-Policy Evaluation with Gumbel-Max Structural Causal Models

Michael Oberst¹ David Sontag¹

Abstract

We introduce an off-policy evaluation procedure for highlighting episodes where applying a reinforcement learned (RL) policy is likely to have produced a substantially different outcome than the observed policy. In particular, we introduce a class of structural causal models (SCMs) for generating counterfactual trajectories in finite partially observable Markov Decision Processes (POMDPs). We see this as a useful procedure for off-policy “debugging” in high-risk settings (e.g., healthcare); by decomposing the expected difference in reward between the RL and observed policy into specific episodes, we can identify episodes where the counterfactual difference in reward is most dramatic. This in turn can be used to facilitate review of specific episodes by domain experts. We demonstrate the utility of this procedure with a synthetic environment of sepsis management.

1. Introduction

When a patient dies in the hospital, we might ask: Could it have been avoided, if the clinicians acted differently? This is impossible to know, because we cannot go back in time. Nonetheless, there is precedent for trying to answer these counterfactual questions: In medical malpractice, for instance, establishing fault requires showing that the injury would not have occurred “but for” the breach in the standard of care (Bal, 2009; Encyclopedia, 2008).

When considering the deployment of new technologies in high-risk environments, such as RL policies in healthcare, we might worry about a similar counterfactual: If we had deployed an RL policy, would it have caused the death of patients who lived under a physician policy? Or more

optimistically, are there lives which could have been saved? This question becomes increasingly relevant, with growing interest in applying RL to everything from HIV therapy (Parbhoo et al., 2017) to epilepsy (Guez et al., 2008) and sepsis (Komorowski et al., 2018).

However, evaluation of RL policies with observational data remains a challenge, albeit a necessary one due to the safety risk of training an RL agent on real patients. Off-policy evaluation in this context is prone to issues such as confounding, small effective sample sizes, and lack of introspection as discussed in Gottesman et al. (2019). Lacking a surefire solution, a physician might review patient trajectories alongside the suggested actions from an RL policy, and ask: Does this seem reasonable? What do I think would have happened if this policy had been followed? Yet, manual inspection of trajectories is not only inefficient, but difficult to interpret without more information: If we are to discover new insights about treatment, shouldn’t there be *some* disagreement?

This motivates our main *conceptual* contribution: By flipping the counterfactual question around, and “asking the algorithm” what would have happened, we aim to highlight the most interesting cases of disagreement, while also demonstrating implicitly why the learned policy is preferred, based on the projected differences in trajectories.

More specifically, in the context of model-based RL with discrete states, we give a post-hoc method to elicit the answers to these questions based on a learned dynamics model. Our goal is to pair counterfactual trajectories with observed trajectories, so that domain experts can “sanity-check” a proposed policy, potentially with additional side-information (e.g., chart review in the case of a patient). In Section 5 we use a synthetic environment of sepsis management to demonstrate how this type of introspection can highlight dangerous implicit assumptions of the learned policy, even when off-policy evaluation is overly optimistic.

Towards generating these counterfactual trajectories, we have to deal with a fundamental issue of non-identifiability in the absence of deterministic dynamics. As we show in this paper, even with a fully-specified finite Markov Decision Process (MDP), there are multiple parameterizations of a structural causal model (SCM) which are equally compatible with the transition and reward distributions, but which sug-

¹CSAIL, Massachusetts Institute of Technology, Cambridge, MA, USA. Correspondence to: Michael Oberst <moberst@mit.edu>.

gest different counterfactual outcomes. For binary treatment and outcomes, the *monotonicity* condition is sufficient to identify the counterfactual distribution (Pearl, 2000) and has been previously used in epidemiology applications (Cuellar & Kennedy, 2018). However, to our knowledge, there is no analogous condition for the categorical case.

This motivates our main *theoretical* contribution, which is two-fold. First, we introduce a general condition of *counterfactual stability* for SCMs with categorical variables (e.g., state transitions) and prove that this condition implies the monotonicity condition in the case of binary categories. Second, we introduce the *Gumbel-Max SCM*, based on the Gumbel-Max trick for sampling from discrete distributions, and (a) demonstrate that it satisfies the counterfactual stability condition, and (b) give a Monte Carlo procedure for drawing counterfactual trajectories under this model. We note that any discrete probability distribution can be sampled using a Gumbel-Max SCM; As a result, drawing counterfactual trajectories can be done in a post-hoc fashion, given any probabilistic model of dynamics with discrete states.

The paper proceeds as follows: In Section 2, we review the formulation of structural causal models and counterfactual estimation, as given by Pearl (2009) and Peters et al. (2017). In Section 3, we introduce the assumption of counterfactual stability, propose a SCM which satisfies it, and show how it can be used to estimate counterfactual trajectories in a Partially Observed Markov Decision Process (POMDP). In Section 4, we discuss related work, and in Section 5 we give an illustrative application using a synthetic example.

2. Preliminaries

2.1. Structural Causal Models

First, we briefly review the concept of *structural causal models*, and encourage the reader to refer to Pearl (2009) (Section 7.1) and Peters et al. (2017) for more details.

Notation: As a general rule throughout, we refer to a random variable with a capital letter (e.g., X), the value it obtains as a lowercase letter (e.g., $X = x$), and a set of random variables with boldface font (e.g., $\mathbf{X} = \{X_1, \dots, X_n\}$). Consistent with Peters et al. (2017) and Buesing et al. (2019), we write P_X for the distribution of a variable X , and p_x for the density function.

Definition 1 (Structural Causal Model (SCM)). A structural causal model \mathcal{M} consists of a set of independent random variables $\mathbf{U} = \{U_1, \dots, U_n\}$ with distribution $P(\mathbf{U})$, a set of functions $\mathbf{F} = \{f_1, \dots, f_n\}$, and a set of variables $\mathbf{X} = \{X_1, \dots, X_n\}$ such that $X_i = f_i(\mathbf{PA}_i, U_i), \forall i$, where $\mathbf{PA}_i \subseteq \mathbf{X} \setminus X_i$ is the subset of \mathbf{X} which are parents of X_i in the causal DAG \mathcal{G} . As a result, the prior distribution $P(\mathbf{U})$ and functions \mathbf{F} determine the distribution $P^{\mathcal{M}}$.




Figure 1. Example translation of a Bayesian Network into the corresponding Structural Causal Model. **Left:** Causal DAG on an outcome Y , covariates X , and treatment T . Given this graph, we can perform do-calculus (Pearl, 2009) to estimate the impact of interventions such as $\mathbb{E}[Y|X, do(T = 1)] - \mathbb{E}[Y|X, do(T = 0)]$, known as the Conditional Average Treatment Effect (CATE). **Right:** All observed random variable are assumed to be generated via structural mechanisms f_x, f_t, f_y via independent latent factors U which cannot be impacted via interventions. Following convention of Buesing et al. (2019), calculated values are given by black boxes (and in this case, are observed), observed latent variables are given in grey, and unobserved variables in white.

As a motivating example to simplify exposition, we will assume the pair of causal graphs given in Figure 1. This causal graph corresponds to a single-action setting which is often assumed in the literature of Conditional Average Treatment Effect (CATE) estimation (see, e.g., Johansson et al., 2016). It can be used to represent, for instance, the effect of a medical treatment T on an outcome Y in the presence of confounding variables X .

The SCM \mathcal{M} defines a complete data-generating processes, which entails the *observational distribution* $P(\mathbf{X}, Y, T)$. It also defines *interventional distributions*. For example, the estimate of CATE is given by $\tau_x = \mathbb{E}[Y|X, do(T = 1)] - \mathbb{E}[Y|X, do(T = 0)]$, where the *do-operator* (Pearl, 2009) is used to signify an intervention.

Definition 2 (Interventional Distribution). Given an SCM \mathcal{M} , an intervention $I = do(X_i := \tilde{f}(\tilde{\mathbf{PA}}_i, \tilde{U}_i))$ corresponds to replacing the structural mechanism $f_i(\mathbf{PA}_i, U_i)$ with $\tilde{f}_i(\tilde{\mathbf{PA}}_i, \tilde{U}_i)$. This includes the concept of atomic interventions, where we may write more simply $do(X_i = x)$. The resulting SCM is denoted \mathcal{M}^I , and the resulting interventional distribution is denoted $P^{\mathcal{M};I}$.

For instance, suppose that Y corresponds to a favorable binary outcome, such as 5-year survival, and T corresponds to a treatment. Then several quantities of interest in causal effect estimation, such as $p^{T=1} := \mathbb{E}[Y|X, do(T = 1)]$, $p^{T=0} := \mathbb{E}[Y|X, do(T = 0)]$, and $\tau_x = p^{T=1} - p^{T=0}$, are all defined by the interventional distribution, which is *forward-looking*, telling us what might be expected to occur if we applied an intervention.

However, we can also define the *counterfactual distribution* which is *retrospective*, telling us what might have happened had we acted differently. For instance, we might ask: Having given the drug and observed that $Y = 1$ (survival), what

would have happened if we had instead withheld the drug?

Definition 3 (Counterfactual Distribution). Given an SCM \mathcal{M} and an observed assignment $\mathbf{X} = \mathbf{x}$ over any set of observed variables, the counterfactual distribution $P^{\mathcal{M}|\mathbf{X}=\mathbf{x};I}$ corresponds to the distribution entailed by the SCM \mathcal{M}^I using the posterior distribution $P(\mathbf{U}|\mathbf{X} = \mathbf{x})$.

Explicitly, given an SCM \mathcal{M} , the counterfactual distribution can be estimated by first inferring the posterior over latent variables, e.g., $P(\mathbf{U}|\mathbf{X} = \mathbf{x}, T = 1, Y = 1)$ in our running example, and then passing that distribution through the structural mechanisms in a modified \mathcal{M}^I (e.g., $I = do(T = 0)$) to obtain a counterfactual distribution over any variable¹.

2.2. (Non)-Identifiability of Counterfactuals

Given an SCM \mathcal{M} , we can compute an answer to our counterfactual question: Having given the drug and observed that $Y = 1$ (survival), what *would have happened* if we had instead withheld the drug? In the binary case, this corresponds to the *Probability of Necessity* (PN) (Pearl, 2009; Dawid et al., 2015), because it represents the probability that the exposure $T = 1$ was necessary for the outcome.

However, the answer to this question is not identifiable without further assumptions: That is, there are multiple SCMs which are all consistent with the interventional distribution, but which produce different counterfactual estimates of quantities like the Probability of Necessity (Pearl, 2009).

2.3. Monotonicity Assumption to Identify Binary Counterfactuals

Nonetheless, there are plausible (though untestable) assumptions we can make that identify counterfactual distributions. In particular, the *monotonicity assumption* (Pearl, 2000; Tian & Pearl, 2000) is sufficient to identify the Probability of Necessity and related quantities used in epidemiology (Cuellar & Kennedy, 2018; Yamada & Kuroki, 2017) to answer counterfactual queries with respect to a binary treatment and outcome.

Definition 4 (Monotonicity). An SCM of a binary variable Y is monotonic relative to a binary variable T if and only if it has the following property^{2,3}: $\mathbb{E}[Y|do(T = t)] \geq \mathbb{E}[Y|do(T = t')] \implies f_y(t, u) \geq f_y(t', u), \forall u$. We can

¹Called abduction, action, and prediction in Pearl (2009)

²We could also write this property as conditional on X

³This definition differs slightly from the presentation of monotonicity in Pearl (2009), where $f_y(t, u)$ being monotonically increasing in t is given as the property, with the testable implication that $\mathbb{E}[Y|do(T = t)] \geq \mathbb{E}[Y|do(T = t')]$ for $t \geq t'$. Because the direction of monotonicity is only compatible with the corresponding direction of the expected interventional outcomes, we fold this into the definition of monotonicity directly, to align with our later definition of counterfactual stability. Also note that we use the notation $Y_{do(T=t)} := f_y(t, u)$ here




Figure 2. SCM for a POMDP as taken from Buesing et al. (2019), with initial state $U_{s1} = S_1$, states S_t , and histories H_t , where the action is generated via the mechanism $\pi(U_a, H_t)$.

write equivalently that the following event never occurs, in the case where $\mathbb{E}[Y|do(T = 1)] \geq \mathbb{E}[Y|do(T = 0)]$: $Y_{do(T=1)} = 0 \wedge Y_{do(T=0)} = 1$. Conversely for $\mathbb{E}[Y|do(T = 1)] \leq \mathbb{E}[Y|do(T = 0)]$, the following event never occurs: $Y_{do(T=1)} = 1 \wedge Y_{do(T=0)} = 0$.

This assumption restricts the class of possible SCMs to those which all yield equivalent counterfactual distributions over Y . For instance, the following SCM exhibits the monotonicity property, and replicates any interventional distribution where $g(x, t) = \mathbb{E}[Y|X = x, do(T = t)]$:

$$Y = \mathbb{1}[U_y \leq g(x, t)], \quad U \sim \text{Unif}(0, 1)$$

More importantly, the monotonicity assumption identifies several counterfactual quantities, such as the Probability of Necessity mentioned earlier.

2.4. Relationship to POMDPs

As noted in Buesing et al. (2019), we can view an episodic Partially Observable Markov Decision Process (POMDP) as an SCM, as shown in Figure 2, where S_t corresponds to states, A_t corresponds to actions, O_t corresponds to observable quantities (including reward R_t), H_t contains history up to time t , i.e., $H_t = \{O_1, A_1, \dots, A_{t-1}, O_t\}$, and stochastic policies are given by $\pi(a_t|h_t)$.

Buesing et al. (2019) investigate a grid-world setting where the transition distribution $P(S_{t+1}|S_t, A_t)$ is deterministic. To extend this counterfactual reasoning to stochastic transitions, we need to make further assumptions on the functional relationships themselves. However, to the best of our knowledge, there is no investigation of plausible assumptions in the literature to identify the answers to counterfactual queries in the case of categorical distributions.

3. Gumbel-Max SCMs

In the context of reinforcement learning with POMDPs, we are typically concerned with estimating the expected reward of a proposed policy $\hat{\pi}$. To formalize notation, a given policy π implies a density over trajectories $\tau \in$

$\mathcal{T} = (S_1, O_1, A_1, \dots, A_{T-1}, S_T, O_T)$, which we denote as $p^\pi(\tau)$, and we let $R(\tau)$ be the total reward of a trajectory τ . For ease of notation, we sometimes write $\mathbb{E}_{\hat{\pi}}$ and \mathbb{E}_{obs} to indicate an expectation taken with respect to $\tau \sim p^{\hat{\pi}}$ and $\tau \sim p^{\pi_{obs}}$ respectively, where $\hat{\pi}$ refers to the proposed (‘target’) policy, and π_{obs} to the observed (‘behavior’) policy.

If we wish to compare the performance of a proposed policy $\hat{\pi}$ and the observed policy π_{obs} , we might compare the difference in expected reward. The expected reward under π_{obs} can be estimated in this case using observed trajectories, without a model of the environment. The difference in expected reward is conceptually similar to the average treatment effect (ATE) of applying the proposed vs. observed policy, and we denote it as δ :

$$\delta := \mathbb{E}_{\hat{\pi}}[R(\tau)] - \mathbb{E}_{obs}[R(\tau)] \quad (1)$$

However, it may be useful to drill down into specific cases: Perhaps there are certain environments, for instance, in which the proposed policy would perform better or worse than the observed policy. One natural decomposition is to condition on the first observed state to estimate a conditional expected reward, e.g.,

$$\delta_o := \mathbb{E}_{\hat{\pi}}[R(\tau)|O_1 = o] - \mathbb{E}_{obs}[R(\tau)|O_1 = o] \quad (2)$$

Equation 2 corresponds conceptually to CATE estimation, where we condition only on pre-treatment information (in this case, O_1 , which occurs before the first action). We can use post-treatment information to decompose this further, over actual trajectories that we have observed, to highlight differences between the observed and proposed policy. Given a causal model of the environment (in the form of an SCM as described in Section 2.4), we can decompose Equation 2 further as follows:

Lemma 1 (Counterfactual Decomposition of Expected Reward). *Let trajectories τ be drawn from $p^{\pi_{obs}}$. Let $\tau_{\hat{\pi}}$ be a counterfactual trajectory, drawn from our posterior distribution over the exogenous U variables under the new policy $\hat{\pi}$. Note that under the SCM, $\tau_{\hat{\pi}}$ is a deterministic function of the exogenous U variables, so we can write $\tau_{\hat{\pi}}(u)$ to be explicit:*

$$\begin{aligned} & \mathbb{E}_{\hat{\pi}}[R(\tau)|O_1 = o] \\ &= \int_{\tau} p^{\pi_{obs}}(\tau|O_1 = o_1) \mathbb{E}_{u \sim p^{\pi_{obs}}(u|\tau)} [R(\tau_{\hat{\pi}}(u))] d\tau \end{aligned}$$

Corollary 1 (Counterfactual Decomposition of δ_o).

$$\begin{aligned} \delta_o &:= \mathbb{E}_{\hat{\pi}}[R(\tau)|O_1 = o] - \mathbb{E}_{obs}[R(\tau)|O_1 = o] \\ &= \int_{\tau} p^{\pi_{obs}}(\tau|O_1 = o_1) \mathbb{E}_{u \sim p^{\pi_{obs}}(u|\tau)} [R(\tau_{\hat{\pi}}(u)) - R(\tau)] d\tau \end{aligned}$$

The proofs of Lemma 1 and Corollary 1 are very similar to (and are essentially implied by) Lemma 1 from Buesing et al. (2019), but are given in self-contained form in the supplement. Corollary 1 implies that we can decompose the expected difference in reward between the policies into differences on *observed episodes* over counterfactual trajectories. However, we face a non-identifiability issue when transitions are not deterministic: Multiple SCMs can all entail the same interventional distribution, but a different set of counterfactual trajectories, and therefore a different decomposition under Lemma 1.

This motivates the theoretical work of this section: We must make our assumptions carefully, as they cannot be tested by data, so it is worth investigating which assumptions are consistent with our causal intuition. We illustrate this non-identifiability (with respect to categorical distributions) in Section 3.1. Then we introduce the condition of *counterfactual stability* (in Section 3.2) for a discrete distribution on k categories, and show that it is compatible with the monotonicity condition of Pearl (2000) in that it implies the monotonicity assumption when $k = 2$. Then we introduce the Gumbel-Max SCM for discrete variables in Section 3.3, and prove that it satisfies the counterfactual stability condition. Finally, in Section 3.4 we describe how the Gumbel-Max SCM assumption can be used to generate a posterior over the counterfactual distribution, under the condition that the data comes from this SCM.

3.1. Non-Identifiability of Categorical Counterfactuals

We will first illustrate that the non-identifiability of counterfactual distributions applies to categorical distributions as well. Consider the categorical distribution over k categories, e.g., the transition kernel $P(S'|S = s, A = a)$ over discrete states. Let $p_i := P(S' = i|S = s, A = a)$. There are multiple ways that we could sample from this distribution with a structural mechanism f and latent variables U . For instance, we could define an ordering **ord** on the categories, and define k intervals of $[0, 1]$ as $[0, p_{\mathbf{ord}(1)}], [p_{\mathbf{ord}(1)}, \sum_{i=1}^2 p_{\mathbf{ord}(i)}], \dots, [\sum_{i=1}^{k-1} p_{\mathbf{ord}(i)}, 1]$. Then we could draw $U \sim \text{Unif}(0, 1)$, and return the interval that u falls into.

However, different permutations **ord** will yield equivalent interventional distributions but can imply different counterfactual distributions. For instance, let $k = 4$ and $p_1 = p_2 = 0.25, p_3 = 0.3, p_4 = 0.2$ and consider an intervention $A = a'$ which defines a different distribution $p'_1 = 0, p'_2 = 0.25, p'_3 = 0.25, p'_4 = 0.5$. Now consider two permutations, **ord** = [1,2,3,4] and **ord'** = [1,2,4,3], and the counterfactual distribution under a' given that $S' = 2, A = a$. In each case, posterior inference over U implies that $P(U|S' = 2, S = s, A = a) \sim \text{Unif}[0.25, 0.5]$. However, under **ord** this implies the counterfactual $S' = 3$,

while under **ord** it implies $S' = 4$. A visual depiction of this can be found in the supplement.

Note that in this example, the mechanism f_{ord} implied a non-intuitive counterfactual outcome: Even though the intervention $A = a'$ lowered the probability of $S' = 3$ (relative to the probability under $A = a$) without modifying the probability of $S' = 2$, it led to a counterfactual outcome of $S' = 3$. Since all choices for **ord** imply the same interventional distribution, there is no way to distinguish between these mechanisms with data.

This motivates the following sections, where we posit a desirable property for categorical SCMs which rules out this result (among others) and is compatible with the notion of monotonicity introduced by Pearl (2000). We then demonstrate that a mechanism based on sampling independent Gumbel variables satisfies this property, which motivates the use of the Gumbel-Max SCM for performing counterfactual inference in this setting.

3.2. Counterfactual Stability

We now introduce our first contribution, the desired property of *counterfactual stability* for categorical SCMs with k categories, laid out in Definition 5. This property would rule out the non-intuitive counterfactual implications of f_{ord} in Section 3.1. We then demonstrate that this condition implies the monotonicity condition when $k = 2$.

Notation: Denote the interventional probability distribution of a categorical variable Y with k categories as $P^{\mathcal{M};I}(Y) = \mathbf{p}$ under intervention I , and \mathbf{p}' under intervention I' , where \mathbf{p}, \mathbf{p}' are vectors in the probability simplex over k categories. To simplify notation for interventional outcomes, we will sometimes denote by Y_I the observed outcome Y under intervention I , and denote by $Y_{I'}$ the counterfactual outcome under intervention I' , such that p_i and $P(Y_I = i)$ are both equivalent to $P^{\mathcal{M};I}(Y = i)$, and similarly for I' . For counterfactual outcomes, we will write $P^{\mathcal{M}|Y_I=i;I'}(Y)$ for the counterfactual distribution of Y under intervention I' given that we observed $Y = i$ under the intervention I .

Definition 5 (Counterfactual Stability). An SCM of a categorical variable Y satisfies *counterfactual stability* if it has the following property: If we observe $Y_I = i$, then for all $j \neq i$, the condition $\frac{p'_i}{p_i} \geq \frac{p'_j}{p_j}$ implies that $P^{\mathcal{M}|Y_I=i;I'}(Y = j) = 0$. That is, if we observed $Y = i$ under intervention I , then the counterfactual outcome under I' cannot be equal to $Y = j$ unless the multiplicative change in p_i is less than the multiplicative change in p_j .

Corollary 2. If \mathcal{M} is a SCM which satisfies counterfactual stability, then if we observe $Y_I = i$, and $\frac{p'_i}{p_i} \geq \frac{p'_j}{p_j}$ holds for all $j \neq i$, then $P^{\mathcal{M}|Y_I=i;I'}(Y = i) = 1$.

This definition and corollary encode the following intuition

about counterfactuals: If we had taken an alternative action that would have *only increased* the probability of $Y = i$, without increasing the likelihood of other outcomes, then the same outcome would have occurred in the counterfactual case. Moreover, in order for the outcome to be different under the counterfactual distribution, the relative likelihood of an alternative outcome must have increased relative to that of the observed outcome. The connection to monotonicity is given in Theorem 1. The proof is straightforward, and found in the supplement.

Theorem 1. Let $Y = f_y(t, u)$ be the SCM for a binary variable Y , where T is also a binary variable. If this SCM satisfies the counterfactual stability property, then it also satisfies the monotonicity property with respect to T .

3.3. Gumbel-Max SCM

Unlike monotonicity with binary outcomes and treatments, the condition of counterfactual stability does not obviously imply any closed-form solution for the counterfactual posterior. Thus, we introduce a specific SCM which satisfies this property, and discuss in Section 3.4 how to sample from the posterior distribution in a straightforward fashion.

We recall the following fact, known as the Gumbel-Max trick (Luce, 1959; Yellott, 1977; Yuille. & L, 2011; Hazan & Jaakkola, 2012; Maddison et al., 2014; Hazan et al., 2016; Maddison et al., 2017):

Definition 6 (Gumbel-Max Trick). We can sample from a categorical distribution with k categories as follows, where \tilde{p}_i is the unnormalized probability $P(Y = i)$: First, draw g_1, \dots, g_k from a standard Gumbel, which can be achieved by drawing u_1, \dots, u_k iid from a $\text{Unif}(0, 1)$, and assigning $g_i = -\log(-\log u_i)$. Then, set the outcome j by taking $\arg \max_j \log \tilde{p}_j + g_j$.

Clearly, we can perform this for any categorical distribution, e.g., the transition distribution $p_i = P(S' = i|S, A)$; In particular, for any discrete variable Y whose parents in a causal DAG are denoted \mathbf{X} , a *Gumbel-Max SCM* assumes the following causal mechanism, where $\mathbf{g} = (g_1, \dots, g_k)$ are independent Gumbel variables:

$$Y = f_y(\mathbf{x}, \mathbf{g}) := \arg \max_j \{\log P(Y = j|\mathbf{X} = \mathbf{x}) + g_j\}$$

Like any mechanism which replicates the conditional distribution under intervention, this mechanism is indistinguishable from any other causal mechanism based on data alone. That said, it does satisfy the property given in Definition 5.

Theorem 2. The Gumbel-Max SCM satisfies the counterfactual stability condition.

The proof is straightforward, and given in the supplement. The intuition is that, when we consider the counterfactual distribution, the Gumbel variables are fixed. Thus, in order

for the argmax (our observed outcome) to change in the counterfactual, the log-likelihood of an alternative outcome must have increased relative to our observed outcome.

3.4. Posterior Inference in the Gumbel-Max SCM

Given a Gumbel-Max SCM as defined above, where $Y = \arg \max_j \log p_j + g_j$ and $p_j := P(Y_I = j)$, we can draw Monte Carlo samples from the posterior $P(\mathbf{g} | Y_I = i)$ using one of two approaches: First, we can use rejection sampling, drawing samples from the prior $P(\mathbf{g})$ and rejecting those where $i \neq \arg \max_j \log p_j + g_j$. Alternatively, it is known (Maddison et al., 2014; Maddison & Tarlow, 2017) that in the posterior, the maximum value and the argmax of the shifted Gumbel variables $\log p_j + g_j$ are independent, and the maximum value is distributed as a standard Gumbel (in the case of normalized probabilities). Thus, we can sample the maximum value first, and then sample the remaining values from shifted Gumbel distributions that are truncated at this maximum value. Then, for each index j , subtracting off the location parameter $\log p_j$ will give us a sample of g_j . We can then add this sample \mathbf{g} to the log-probabilities under I' (i.e., $\log \mathbf{p}'$) and take the new argmax to get a sample of the counterfactual outcome Y under intervention I' .

4. Related Work

In the machine learning community, there has been recent interest in applying insights from causal inference to augment the training of RL models, such as in bandit settings (Lee & Bareinboim, 2018) and in model-based RL (Buesing et al., 2019). The most similar work to our own is Buesing et al. (2019), but while they use counterfactuals to approximate draws from the interventional distribution, we treat the counterfactual distribution as the primary object of interest.

The term ‘counterfactuals’ is often used more broadly in causal inference and machine learning, e.g., in the estimation of potential outcomes (Imbens & Rubin, 2015), which are also counterfactual outcomes at an individual level. However, this is often used (e.g., Schulam & Saria, 2017; Johansson et al., 2016) primarily towards predicting quantities related to the interventional distribution, such as CATE.

Assumptions on structural mechanisms are also used in structure learning to identify the direction of causation $X \rightarrow Y$ or $X \leftarrow Y$ from the observational distribution alone, with both continuous (Peters & Schölkopf, 2014; Mooij et al., 2016) and discrete (Kocaoglu et al., 2017) variables. However, these assumptions imply differences in the observational distribution (by design), whereas our assumptions distinguish between counterfactual distributions which are both observationally and interventionally equivalent.

The assumption of monotonicity implicitly appears in early work in epidemiology on estimating quantities like the ‘rel-

ative risk ratio’ (Miettinen, 1974), which are often imbued with causal interpretations (Pearl, 2009; Yamada & Kuroki, 2017). Formalizing the assumption of monotonicity, required to correctly impute causal meaning to these quantities, is covered in Balke & Pearl (1994); Pearl (2000); Tian & Pearl (2000). More recent work in epidemiology uses the assumption of monotonicity explicitly, e.g., to estimate the counterfactual effect of water sanitation in Kenya (Cuellar & Kennedy, 2018), and there has been ample discussion and debate regarding how this reasoning could apply (in principle) to legal cases, such as litigation around the toxic effects of drugs (Dawid et al., 2016).

The use of Gumbel variables to sample from a categorical distribution (the Gumbel-Max Trick) is well known (see Definition 6) but in recent work in the machine learning community (Maddison et al., 2017; Hazan et al., 2016) it is not imbued with the same causal / counterfactual interpretation that we propose here. The Gumbel-Max mechanism was initially introduced in the discrete-choice literature (Luce, 1959), where it is used as a generative model for decision-making under utility maximization (Train, 2002; Aguirregabiria & Mira, 2010), where the log probabilities may be assumed to follow some functional form, such as being linear in features. This is motivated by understanding the impact of different characteristics on consumer choices, see (Aguirregabiria & Mira, 2010, Example 1). In contrast, we decouple this structural mechanism (for estimation of counterfactuals) from the statistical model used to estimate the conditional probability distributions under interventions.

5. Experiments

As discussed in Section 1, our hope is to provide a method for qualitative introspection and ‘debugging’ of RL models, in settings where a domain expert could plausibly examine individual trajectories. We give an illustrative example of this use case here, motivated by recent work examining the use of RL algorithms for treating sepsis among intensive-care unit (ICU) patients (Komorowski et al., 2018). In particular, we use a simple simulator of sepsis and “debug” an RL policy that is learned on observed trajectories.

An analysis like this requires three ingredients: First, we are given *observed trajectories*, but cannot directly interact with the environment⁴. Second, we have access to a *structural causal model* of the environment. In this case, that model is a finite MDP, learned based on observed samples, combined with the assumption of a Gumbel-Max SCM for transition distributions. Finally, we need a *target policy* to evaluate. We refer to the policy which generated the data as the *behavior policy*, to distinguish it from the target policy.

⁴We do not assume access to a simulator; In this example, it is used only for obtaining the initial observed trajectories.

In Sections 5.1-5.2 we describe our illustrative scenario, in which a target RL policy appears to perform well using off-policy evaluation methods such as weighted importance sampling, when it is actually much worse than the behavior policy. In Sections 5.3-5.4 we then demonstrate how our method could be used to identify a promising subset of trajectories for further introspection, and uncover the flaws in the target policy using side information (e.g., chart review of individual patients). All the code required to reproduce our experiments is available online at <https://www.github.com/clinicalml/gumbel-max-scm>

5.1. Illustrative Analysis: Sepsis

Environment: Our simulator includes four vital signs (heart rate, blood pressure, oxygen concentration, and glucose levels) with discrete states (e.g., low, normal, high), along with three treatment options (antibiotics, vasopressors, and mechanical ventilation), all of which can be applied at each time step. Reward is +1 for discharge of a patient, and -1 for death. Discharge occurs only when all patient vitals are within normal ranges, and all treatments have been stopped. Death occurs if at least three of the vital signs are simultaneously out of the normal range. In addition, a binary variable for diabetes is present with 20% probability, which increases the likelihood of fluctuating glucose levels.

Observed Trajectories: For the purposes of this illustration, the behaviour policy was constructed using Policy Iteration (Sutton & Barto, 2017) with full access to the parameters of the underlying MDP (including diabetes state). This was done deliberately to set up a situation in which the observed policy performs well. To introduce variation, the policy takes a random alternative action w.p. 0.05. Using this policy, we draw 1000 patient trajectories from the simulator, with a maximum of 20 time steps. If neither death nor discharge is observed, the observed reward is zero.

Structural Causal Model: For this illustration, we ‘hide’ glucose and diabetes state in the observed trajectories; Given this reduced state-space, we learn the parameters of the finite MDP by using empirical counts of transitions and rewards from the 1000 observed trajectories, with death and discharge treated as absorbing states. For state / action pairs that are not observed, we assume that any action leads to death, and confirm that this results in a target policy which never takes an action that has never been observed. For counterfactual evaluation, we make the assumption that the transitions are generated by a Gumbel-Max SCM.

Target Policy: The target policy is learned using Policy Iteration on the parameters of the learned MDP. Because the target policy is learned using a limited number of samples, as well as an incomplete set of variables, it should perform poorly relative to the behavior policy.




Figure 3. Estimated reward under the target (RL) policy, with 95% uncertainty intervals generated through 100 bootstrapped samples (with replacement) of the same 1000 observed trajectories (for 1,2,4) and of 1000 new trajectories under the target policy (for 3,5). (1) **Obs**: Observed reward under the behavior policy. (2) **WIS**: Estimated reward under the target policy using weighted importance sampling. (3) **MB**: Estimated reward using the learned MDP as a generative model. (4) **CF**: Estimated reward over counterfactual trajectories (5 per observed trajectory). (5) **True**: Observed reward under the target policy, over 1000 newly simulated trajectories.

5.2. Off-Policy Evaluation Can Be Misleading

First, we demonstrate what might be done to evaluate this target policy without the use of counterfactual tools. In Figure 3, we compare the observed reward of the actual trajectories against the estimated reward of the target policy. Using weighted importance sampling on the given trajectories, the target policy appears superior to the behavior policy. We also use the parameters of the learned MDP to perform model-based off-policy evaluation (MB-PE), using the MDP as a generative model to simulate trajectories and their expected reward. Both of these suggest that the target policy is superior to the behavior policy. In reality, the target policy is inferior (as expected by construction), as verified by drawing new samples from the simulator under the target policy. This corresponds conceptually to what would happen if the target policy were deployed “in the wild”.

With this in mind, we demonstrate how examining individual counterfactual trajectories gives insight into the target policy. The first step is to apply counterfactual off-policy evaluation (CF-PE) using the same MDP and the Gumbel-Max SCM. This yields similarly optimistic results as MB-PE. However, by pairing counterfactual outcomes with observed outcomes of individual patients, we can investigate *why* the learned MDP concludes (wrongly) that the target policy would be so successful.

5.3. Identification of Informative Trajectories

To debug this model (without access to a simulator), we can start by drawing counterfactual trajectories for each individual patient under the target policy. With these in hand, we can assign each individual patient to one of nine

Observed Outcome	Died	1%	0%	10%
	Disch. No Chg.	0%	1%	43%
	Died No Chg. Disch.	0%	0%	44%
		Counterfactual Outcome		

Figure 4. Decomposition of 1000 observed patient trajectories based on observed outcome (Died, no change, and discharged) vs counterfactual outcome under the target policy, using the most common outcome over 5 draws from the counterfactual posterior.

categories, based on the most frequently occurring counterfactual outcome (death, no change, or discharge) in Figure 4. This highlights individual trajectories for further analysis, as discussed in the next section⁵.

5.4. Insights from Examining Individual Trajectories

Using this decomposition, we can focus on the 10% of observed trajectories where the model concludes that “if the physician had applied the target policy, these patients who died would have most likely lived”.

This is a bold statement, but also one that is plausible for domain experts to investigate (e.g., through chart review of these specific patients), to try and understand the rationale. We illustrate this type of analysis in Figure 5, which shows both the observed trajectory and the counterfactual trajectories for a simulated patient.

This example illustrates a dangerous failure mode, where the target policy would have halted treatment despite the glucose vital being dangerously low (e.g., at $t = 5, 7, 8, 11$). Under the learned MDP, the counterfactual optimistically shows a speedy discharge as a result of halting treatment. To understand why, recall that discharge occurs when all four vitals are normal and treatment is stopped. Because diabetes and glucose fluctuations are relatively rare, and because the MDP does not observe either, the model learns that there is a high probability of discharge when the first three vitals are normal, and the action of ‘stop all treatments’ is applied.

6. Conclusion

Given the desire to deploy RL policies in high-risk settings (e.g., healthcare), it is important to develop more tools and

⁵We only draw 5 counterfactuals per observed trajectory for illustrative purposes here, but note that standard concentration arguments could be used to quantify how many of these independent draws are required to achieve a desired precision on counterfactual quantities of interest, e.g., the probability of death




Figure 5. Observed and counterfactual trajectories of a patient. The first four plots show the progression of vital signs, and the last three show the treatment applied. For vital signs, the normal range is indicated by red dotted lines. The black lines show the observed trajectory, which ends in death (signified by the red cross), and the blue lines show five counterfactual trajectories all of which end in discharge, signified by green diamonds. The glucose vital sign was not included in the model, and hence does not have a counterfactual trajectory. Note how this differs from a newly simulated trajectory of a patient with the same initial state, e.g., all the counterfactual trajectories are identical to the observed trajectory up to a divergence in actions ($t = 2$).

techniques to introspect these models and the policies they learn. Our proposed technique can be used to flag individual trajectories for review by a domain expert, based on the fact that they have divergent counterfactual and observed outcomes. Specifically, we gave an example of how this could help to identify sub-optimal (and perhaps dangerous) behavior in a proposed policy.

Applying this approach requires knowing the structural causal model: simply knowing the MDP is insufficient, as demonstrated by our unidentifiability results. The Gumbel-Max SCM is an example of an SCM that may be realistic in some settings. As followup, it would be interesting to understand the sensitivity of counterfactual trajectories to the specific choice of SCM. For instance, while we are not aware of another SCM that both satisfies the counterfactual stability condition and can model any discrete conditional probability distribution, it would be interesting to characterize the space of such alternative SCMs, or prove that no other such SCM exists for $k > 2$ which satisfies these conditions.

Acknowledgments

We thank Christina X. Ji and Fredrik D. Johansson for their work on developing an earlier version of the sepsis simulator. We also thank Fredrik D. Johansson for insightful comments and feedback, as well as Alex D’Amour, Uri Shalit, and David Blei for helpful discussions. This work was supported in part by Office of Naval Research Award No. N00014-17-1-2791, and Michael was supported in part by the Analog Devices Graduate Fellowship.

References

- Aguirregabiria, V. and Mira, P. Dynamic discrete choice structural models: A survey. *Journal of Econometrics*, 156(1):38–67, 2010. ISSN 03044076. doi: 10.1016/j.jeconom.2009.09.007. URL <http://dx.doi.org/10.1016/j.jeconom.2009.09.007>.
- Bal, B. S. An introduction to medical malpractice in the United States. *Clinical Orthopaedics and Related Research*, 467(2):339–347, 2009. ISSN 0009921X. doi: 10.1007/s11999-008-0636-2.
- Balke, A. and Pearl, J. Counterfactual Probabilities: Computational Methods, Bounds and Applications. *Proceedings of the Tenth Conference Annual Conference on Uncertainty in Artificial Intelligence (UAI-94)*, pp. 46–54, 1994. doi: 10.2202/1557-4679.1322.
- Buesing, L., Weber, T., Zwols, Y., Heess, N., Racaniere, S., Guez, A., and Lespiau, J.-B. Woulda, Coulda, Shoulda: Counterfactually-Guided Policy Search. In *International Conference on Learning Representations*, 2019. URL <https://openreview.net/forum?id=BJG0voC9YQ>.
- Cuellar, M. and Kennedy, E. H. A nonparametric projection-based estimator for the probability of causation, with application to water sanitation in Kenya. *Journal of the Royal Statistical Society: Series A (Special Issue on Causal Inference)*, pp. 1–24, 2018.
- Dawid, P., Faigman, D. L., and Fienberg, S. E. On the Causes of Effects: Response to Pearl. *Sociological Methods and Research*, 44(1):165–174, 2015. ISSN 15528294. doi: 10.1177/0049124114562613.
- Dawid, P., Musio, M., and Fienberg, S. E. From statistical evidence to evidence of causality. *Bayesian Analysis*, 11(3):725–752, 2016. ISSN 19316690. doi: 10.1214/15-BA968.
- Encyclopedia, W. *West’s Encyclopedia of American Law*. The Gale Group, 2nd edition, 2008. URL <https://legal-dictionary.thefreedictionary.com/proximate+cause>.
- Gottesman, O., Johansson, F., Komorowski, M., Faisal, A., Sontag, D., Doshi-Velez, F., and Celi, L. A. Guidelines for reinforcement learning in healthcare. *Nature Medicine*, 25(1):16–18, 2019. ISSN 1078-8956. doi: 10.1038/s41591-018-0310-5. URL <http://www.nature.com/articles/s41591-018-0310-5>.
- Guez, A., Vincent, R. D., Avoli, M., and Pineau, J. Adaptive Treatment of Epilepsy via Batch-mode Reinforcement Learning. *AAAI*, 2008. URL <http://www.cs.mcgill.ca/~jpineau/files/guez-iaai08.pdf>.
- Hazan, T. and Jaakkola, T. On the partition function and random maximum a-posteriori perturbations. In *ICML*, 2012.
- Hazan, T., Papandreou, G., and Tarlow, D. *Perturbation, Optimization, and Statistics*. MIT Press, 2016.
- Imbens, G. W. and Rubin, D. B. *Causal Inference for Statistics, Social, and Biomedical Sciences*. Cambridge University Press, 2015. ISBN 978-0-521-88588-1.
- Johansson, F. D., Shalit, U., and Sontag, D. Learning Representations for Counterfactual Inference. In *ICML*, 2016. URL <http://arxiv.org/abs/1605.03661>.
- Kocaoglu, M., Dimakis, A. G., Vishwanath, S., and Hassibi, B. Entropic Causal Inference. In *AAAI*, pp. 1156–1162, 2017. URL <https://arxiv.org/pdf/1611.04035.pdf>.
- Komorowski, M., Celi, L. A., Badawi, O., Gordon, A. C., and Faisal, A. A. The Artificial Intelligence Clinician learns optimal treatment strategies for sepsis in intensive care. *Nature Medicine*, 24(11):1716–1720, 2018. ISSN 1546170X. doi: 10.1038/s41591-018-0213-5. URL <http://dx.doi.org/10.1038/s41591-018-0213-5>.
- Lee, S. and Bareinboim, E. Structural Causal Bandits: Where to Intervene? In Bengio, S., Wallach, H., Larochelle, H., Grauman, K., Cesa-Bianchi, N., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems 31*, pp. 2572–2582. Curran Associates, Inc., 2018.
- Luce, R. D. *Individual Choice Behavior: A Theoretical Analysis*. Wiley, New York, 1959.
- Maddison, C. J. and Tarlow, D. Gumbel Machinery, 2017. URL <https://cmaddis.github.io/gumbel-machinery>.
- Maddison, C. J., Tarlow, D., and Minka, T. A* Sampling. *Advances in Neural Information Processing Systems*, 2014. ISSN 02767783. doi: 10.1016/

- B978-0-12-803459-0.00005-4. URL <http://arxiv.org/abs/1411.0030>.
- Maddison, C. J., Mnih, A., and Teh, Y. W. The Concrete Distribution: A Continuous Relaxation of Discrete Random Variables. In *International Conference on Learning Representations (ICLR)*, 2017. ISBN 0780365402. URL <http://arxiv.org/abs/1611.00712>.
- Miettinen, O. S. Proportion of disease caused or prevented by a given exposure, trait or intervention. *American Journal of Epidemiology*, 99(5):325–332, 1974. ISSN 00029262. doi: 10.1093/oxfordjournals.aje.a121617.
- Mooij, J. M., Peters, J., Janzing, D., Zscheischler, J., and Schölkopf, B. Distinguishing Cause from Effect Using Observational Data: Methods and Benchmarks. *Journal of Machine Learning Research*, 17: 1–102, 2016. URL <http://jmlr.org/papers/volume17/14-518/14-518.pdf>.
- Parbhoo, S., Bogojeska, J., Zazzi, M., Roth, V., and Doshi-Velez, F. Combining Kernel and Model Based Learning for HIV Therapy Selection. *AMIA Joint Summits on Translational Science Proc*, 2017:239–248, 2017. ISSN 2153-4063. URL <http://www.ncbi.nlm.nih.gov/pubmed/28815137><http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=PMC5543338>.
- Pearl, J. Probabilities of causation: three counterfactual interpretations and their identification. *Synthese*, 121(1):93–149, 2000. ISSN 00397857. doi: 10.1023/A:1005233831499.
- Pearl, J. *Causality: Models, Reasoning, and Inference*. Cambridge University Press, 2nd edition, 2009. ISBN 052189560X. URL <https://dl.acm.org/citation.cfm?id=1642718>.
- Peters, J. and Schölkopf, B. Causal Discovery with Continuous Additive Noise Models. *Journal of Machine Learning Research*, 15(June):2009–2053, 2014. ISSN 15337928. doi: 10.1.1.144.4921. URL <http://jmlr.org/papers/volume15/peters14a/peters14a.pdf>.
- Peters, J., Janzing, D., and Schölkopf, B. *Elements of Causal Inference: Foundations and Learning Algorithms*. MIT Press, Cambridge, MA, 2017. ISBN 9780262037310.
- Schulam, P. and Saria, S. Reliable Decision Support using Counterfactual Models. In *31st Conference on Neural Information Processing Systems*, 2017.
- Sutton, R. S. and Barto, A. G. *Reinforcement Learning: An Introduction*. MIT Press, 2nd edition, 2017.
- Tian, J. and Pearl, J. Probabilities of causation : Bounds and identification. *Annals of Mathematics and Artificial Intelligence*, 28(1-4):287–313, 2000. ISSN 1012-2443. doi: 10.1023/A:1018912507879.
- Train, K. *Discrete choice methods with simulation*. Cambridge University Press, 2002.
- Yamada, K. and Kuroki, M. Counterfactual-Based Prevented and Preventable Proportions. *Journal of Causal Inference*, 5(2), 2017. ISSN 2193-3677. doi: 10.1515/jci-2016-0020.
- Yellott, J. I. The relationship between Luce’s choice axiom, Thurstone’s theory of comparative judgment, and the double exponential distribution. *Journal of Mathematical Psychology*, 15(2):109–144, 1977.
- Yuille., G. P. and L, A. Perturb-and-map random fields: Using discrete optimization to learn and sample from energy models. In *ICCV*, 2011.