
Abstract Counterfactuals for Language Model Agents

Edoardo Pona

King's College in London
edoardo.1.pona@kcl.ac.uk

Milad Kazemi

King's College in London
milad.kazemi@kcl.ac.uk

Yali Du

King's College in London
yali.du@kcl.ac.uk

David Watson

King's College in London
david.watson@kcl.ac.uk

Nicola Paoletti

King's College in London
nicola.paoletti@kcl.ac.uk

Abstract

Counterfactual inference is a powerful tool for analysing and evaluating autonomous agents, but its application to language model (LM) agents remains challenging. Existing work on counterfactuals in LMs has primarily focused on token-level counterfactuals, which are often inadequate for LM agents due to their open-ended action spaces. Unlike traditional agents with fixed, clearly defined action spaces, the actions of LM agents are often implicit in the strings they output, making their action spaces difficult to define and interpret. Furthermore, the meanings of individual tokens can shift depending on the context, adding complexity to token-level reasoning and sometimes leading to biased or meaningless counterfactuals. We introduce *Abstract Counterfactuals*, a framework that emphasises high-level characteristics of actions and interactions within an environment, enabling counterfactual reasoning tailored to user-relevant features. Our experiments demonstrate that the approach produces consistent and meaningful counterfactuals while minimising the undesired side effects of token-level methods. We conduct experiments on text-based games and counterfactual text generation, while considering both token-level and latent-space interventions.

1 Introduction

The recent successes of Large Language Models (LLMs) have paved the way for a novel approach to developing autonomous agents. Previously, agents were typically trained in isolated environments with limited knowledge. Language Model (LM) Agents [37] instead leverage their vast background knowledge (owing to their training on internet-scale datasets) to solve increasingly general tasks, including web browsing and research [38], multi-modal robotics [4], and navigating open-ended environments [36, 39]. As such, LM Agents have also been studied for high-risk domains such as medicine, law, and diplomacy [17, 20, 29, 32, 33]. This, however, has important safety implications due to the (well-known) issues surrounding LLM safety, such as social biases and opaque reasoning [1, 13, 21].

Causal and counterfactual explanations are effective techniques to enhance the explainability and reliability of AI models. These techniques can answer how a model would have behaved in alternative (counterfactual) settings, given observations of its behaviour in a factual setting [15, 24, 34]. This ability to reason about “what if” scenarios is crucial in understanding responsibility and blame in autonomous systems [6, 14] and deriving counterfactual policies—i.e., policies which, in hindsight, would have been optimal with minimal interventions [2, 16, 19]. Recently, Ravfogel et al. [27] and Chatzi et al. [7] have proposed methods for counterfactual inference on LLMs based on *structural causal models* (SCMs) [24]. These two methods are the first to apply SCMs for LLM counterfactuals.




Figure 1: Abstract Counterfactuals overview. Our method considers the meaning of the observed action—encoded in the abstraction Y —and in the counterfactual setting, updates actions' probabilities according to the observed value of Y rather than the observed token A .

They define a *token-level* SCM, that models the sampling of individual tokens. We call these approaches *token-level counterfactual (TLCF)*.

We argue that applying token-level inference to LM agents is inadequate for two main reasons. In *open-text environments*, where agents issue actions as arbitrary strings, the lack of a predefined action space requires the environment to interpret language model (LM) outputs. Here, token-level inference may be inadequate for capturing the high-level semantics, emphasizing token-level syntax instead. Additionally, the high-level meaning of actions can vary with the agent's context; identical symbols may correspond to different actions across contexts, particularly in *choice-based* environments where actions are selected from predefined strings. For example, the meaning of tokens such as “choice 1” can differ widely across situations, an issue token-level methods fail to address (as illustrated in the subsequent example). In general, when LM agents serve as ‘algorithmic models’ [3] of real-world systems—for example, in agent-based simulations such as Vezhnevets et al. [35]—it may be preferable to compute counterfactuals at a higher conceptual level, abstracting away from the low-level (token-level) details of their computational implementations. We may also wish to leverage expert knowledge to explicitly define or inform this conceptual abstraction level. This challenge arises not because the token-generation function f_A fails to encode context, but because counterfactual inference at the token level conditions on surface tokens whose meaning may shift or vanish under a changed context.

In this paper, we introduce an approach to LM agents' counterfactuals that overcomes the above limitations of token-level methods. The key idea is to introduce an abstraction Y of the LLM action A , so that $Y | A$ captures the high-level semantics of A and remains valid across contexts. Then, to derive a counterfactual action a' , instead of performing counterfactual inference directly on the tokens of A , our method does it at the level of Y . The resulting counterfactual abstraction Y' is then “mapped back” into the action space, i.e., we derive a' so that its abstraction $Y | a'$ is compatible with Y' . This way, we obtain the desired counterfactual (a') but override the token-level mechanism that generates the action. This also means our method requires only black-box access to the LLM.

To better understand our approach, consider the hypothetical text-based adventure game illustrated in Figure 1. In the factual setting, the LM agent's policy assigns a distribution $P(A | s)$ over the possible action choices (i.e., the game options), encoded by tokens ‘1’, ‘2’, and ‘3’, given the context $S = s$. Suppose the agent samples action ‘2’, which, in s , corresponds to hiding after seeing a bear in the woods, instead of running away (action ‘3’) or facing the bear (action ‘1’). Now consider a counterfactual context s' where a lion in the jungle is nearing the agent. Having observed our agent hiding from the bear in the woods, we want to predict what they would do in the jungle.

If we follow the token-level approaches of [26] and [7] counterfactual inference would increase the counterfactual probability of action 2 (the action token previously observed). The problem is that action ‘2’ in context s' means the agent will face the lion, while in s , the same action token indicates

cautious behaviour. A sensible counterfactual should rather increase the probability of cautious behaviour in s' .

In this example, our technique would introduce an abstraction Y to represent the high-level semantics of the actions. For instance, Y can be a binary variable that represents courage. So, in the factual context s , action ‘2’ maps to $Y = 0$ (not courageous); in the counterfactual context s' , action ‘2’ maps to $Y = 1$ (courageous) while actions ‘1’ and ‘3’ map to $Y = 0$. In our inference approach, mediated by the abstraction Y , observing $Y = 0$ leads to increasing the likelihood of $Y' = 0$ in the counterfactual world, and with that, the likelihoods of the action(s) compatible with that abstraction value (the cautious actions ‘1’ and ‘3’, in this case).

In summary, we introduce *abstract counterfactuals (ACF)*, an approach for LM agents’ counterfactuals that overcomes the issues of token-level action generation mechanisms by leveraging a semantics- and context-aware proxy of the actions. Unlike TLCF approaches, our inference method crucially does not require white-box access to the LLM internals. Moreover, our abstractions can be defined either in an unsupervised fashion, with an auxiliary LM discovering and classifying the relevant abstractions, or through supervised classifiers informed by expert-defined categories.

We evaluate our approach on three benchmarks: MACHIAVELLI [23],¹ a choice-based game for evaluating agents’ social decision making, and two open-text tasks, involving the generation of short biographies [8]² and Reddit comments [9],³ respectively. Results demonstrate that our abstract counterfactuals consistently outperform token-level ones by improving the semantic consistency between counterfactual and factual actions and also within multiple counterfactual realizations.

2 Preliminaries

Structural causal models and counterfactuals [24] *Structural causal models (SCM)* provide a mathematical framework for causal inference. An SCM $\mathcal{C} = (\mathbf{S}, \mathbf{U}, P_{\mathbf{U}})$ consists of a set $\mathbf{S}$ of (acyclic) structural assignments of the form $X_i = f_i(\mathbf{PA}_i, U_i)$ for $1 \leq i \leq |S|$, where X_i are the endogenous (observed) variables, $\mathbf{PA}_i \subseteq \{X_1, \dots, X_{|S|}\} \setminus \{X_i\}$ denote the parents of variable X_i , and $P_{\mathbf{U}} = P_{U_1} \times \dots \times P_{U_{|S|}}$ is the joint distribution over the so-called exogenous (unobservable) variables $\mathbf{U}$. The exogenous distribution $P_{\mathbf{U}}$ and the assignments $\mathbf{S}$ induce a unique distribution over the endogenous variables $\mathbf{X}$, denoted $P_{\mathbf{X}}^{\mathcal{C}}$.

An *intervention* on $\mathcal{C}$ corresponds to replacing one or more structural assignments, thereby obtaining a new SCM $\tilde{\mathcal{C}}$. The entailed distribution of this new SCM, $P_{\mathbf{X}}^{\tilde{\mathcal{C}}}$, is called *interventional distribution* and allows us to predict the causal effect of the intervention on the SCM variables.

Given an observation $\mathbf{x} \sim P_{\mathbf{X}}^{\mathcal{C}}$, *counterfactual inference* corresponds to predicting the hypothetical value of $\mathbf{x}$ had we applied an intervention on $\mathcal{C}$. This process consists in inferring the value of $\mathbf{U}$ that led to $\mathbf{x}$ (known as abduction step), by deriving the posterior

$$P_{\mathbf{U}|\mathbf{X}=\mathbf{x}}(\mathbf{u}) = \frac{P(\mathbf{X} = \mathbf{x} \mid \mathbf{U} = \mathbf{u})P_{\mathbf{U}}(\mathbf{u})}{P(\mathbf{X} = \mathbf{x})}$$

Then, counterfactual statements correspond to statements evaluated on the SCM obtained from $\mathcal{C}$ after performing the target intervention and updating $P_{\mathbf{U}}$ with the posterior $P_{\mathbf{U}|\mathbf{X}=\mathbf{x}}$.

Structural causal models for auto-regressive token generation To perform causal inference in the context of language modelling, Chatzi et al. [7], Ravfogel et al. [27] have recently proposed an SCM to describe the process of auto-regressive token generation, summarised below.

Let V be the LM vocabulary (set of available tokens). Given a token sequence (or *prompt*) $\mathbf{x} \in \bigcup_{j=1}^K V^j$, where K denotes the maximum sequence length, the next token generated X^k by the LM follows a categorical distribution $X^k \sim \text{Cat}(\text{softmax}(\lambda(\mathbf{x}^{k-1})))$, where λ refers to a decoder (e.g., a transformer) which outputs logits over V given an input sequence. To obtain a structural assignment for the categorical variable X^k , Chatzi et al. [7] and Ravfogel et al. [27] employ the Gumbel-Max

¹code MIT; data research-only

²MIT

³CC BY 4.0

SCM [22], so that X^k can be expressed as a deterministic function of the prompt $\mathbf{x}^{k-1}$, the language model λ , and the exogenous noise U :

$$f_{X^k}(\mathbf{x}^{k-1}, \lambda, U) = \arg \max_{v \in V} (\lambda(\mathbf{x}^{k-1})_v + U_v), \quad (1)$$

where $U = (U_v)_{v \in V}$ is a vector of i.i.d. Gumbel random variables and $\lambda(\mathbf{x})_v$ denote the logits output by the language model for token v given the prompt $\mathbf{x}$. At each time step of the sequence generation, the sampled token x^k is appended to the current sequence. This procedure continues until sampling of an ‘end-of-sequence’ token, or the sequence exceeds the maximum length K .

Given a sampled (factual) token x^k , i.e., a realisation of the distribution entailed by the SCM (1), a *token-level counterfactual (TLCF)* is derived from (1) as described in the previous section, where the abduction step corresponds to inferring the Gumbel noise posterior $U' = U \mid x^k, \mathbf{x}^{k-1}$ using the observation x^k and the prompt $\mathbf{x}^{k-1}$. Since the mechanism (1) is non-invertible, U' cannot be uniquely identified and so requires approximate inference [18], resulting in a stochastic U' and a stochastic counterfactual. Interventions may involve altering the LM decoder λ , e.g., by changing its weights, or by manipulating the tokens in the prompt $\mathbf{x}^{k-1}$.

It is essential to notice that *TLCFs always increase the counterfactual probability of the observed token x^k* . Indeed, if x^k was observed, the posterior Gumbel U' will increase the probability of those noise values u where u_{x^k} is high enough to maximize $\lambda(\mathbf{x}^{k-1})_{x^k} + u_{x^k}$ across all possible tokens; if instead u_{x^k} is not high enough to yield x^k , then its posterior probability will be zero. We argue that this property of TLCFs represents a limitation that hinders the semantic consistency of LM agents counterfactuals, as we explain next.

3 Abstract Counterfactuals Method

We now present an SCM for an LM Agent operating in a sequential decision-making environment, similar to the SCMs proposed by [5, 22]. The *state* at step t is given by a pair $S_t = (\mathbf{X}_t, \theta)$, where $\mathbf{X}_t$ represents the agent’s current prompt and θ represents the LM’s parameters (e.g. weights, sampling strategy, latent manipulations)⁴. The agent uses a stochastic, state-conditional policy to select an *action* A_t . After deploying the action, the environment transitions into a new state S_{t+1} . This process is summarised by the following SCM:

$$A_t = f_A(S_t, U_t^A); S_{t+1} = f_S(S_t, A_t, U_t^S), \quad (2)$$

where U_t^A and U_t^S are the exogenous factors associated to the action and state-update mechanisms.

In an LM agent, the action A_t consists of a sequence of tokens sampled from the model, i.e., it is the result of applying the autoregressive token generation SCM (1) for multiple steps. In this case, a TLCF approach would perform posterior inference of the sequence of Gumbel exogenous variables U_t^A given the observed token sequence a_t and the state s_t . Hence, as discussed previously, TLCF would assign a higher counterfactual probability to the observed tokens, irrespective of whether the tokens remain relevant in the counterfactual context. This behaviour leads to two different categories of *failure cases* for TLCF:

- *Choice-based environments*: in such environments, the token-level representation of the action may change its meaning across different contexts; this was discussed earlier in the Figure 1 example, where the observed action token ‘2’ entails a cautious action in the factual context and a reckless action in the counterfactual one. Another example of this failure case is given in Figure 3 of Section 4.1.
- *Open-text environments*: in this case, the action space consists of arbitrary token sequences. TLCF approaches would ignore the high-level meaning of the generated action text, focusing instead on token-level utterances. The result is that the inference procedure would carry no or very little semantic information from the factual/observed context into the counterfactual one. An example of this issue is shown in Appendix I, where, after a gender-steering intervention on the model, TLCF fails to generate a short bio which is consistent with the profession observed in the factual setting.

⁴We treat θ as part of the state because we may wish to intervene on it just as we would intervene on any other variable in the SCM.

In other words, we cannot trust the token-level mechanism f_A alone to perform counterfactual inference (as done instead in previous approaches [7, 27]); quoting [10], we do not want to “*conflate the uncertainty of the [language] model over the meaning of its answer with that over the exact tokens used to express that meaning*”.

Our *abstract counterfactuals* (ACF) method relies on a simple yet effective idea: introduce in the above LM agent SCM (2) an abstraction variable Y_t that represents the high-level meaning of action A_t in state S_t , in a way that the meaning expressed in Y_t remains consistent across contexts. This is illustrated in Figure 2 and described by the following structural assignment:

$$Y_t = f_Y(A_t, S_t, U_t^Y).$$

We stress that Y_t depends on both state and action, allowing us to capture the action’s context-dependent meaning. Our abstraction acts as a proxy for the token-level action A_t , a proxy which retains those high-level features that matter when reasoning about counterfactual outcomes. For instance, in the Figure 1 example, Y describes whether the agent is cautious or courageous; in our experiments with the MACHIAVELLI benchmark (Section 4.1), Y is used to classify the agent’s ethical behaviour, e.g., its intention to cause physical harm; or, Y can represent the notion of ‘job’ in a biography-generation task (see Section 4.2).

Supervised and unsupervised abstractions. The ACF method leaves us the choice of how to construct this abstraction. Where the user has ‘expert knowledge’ about the domain or an opinionated view on what features of the samples should be considered meaningful, we can use a *supervised* abstraction. That is, either using annotations directly (as we do in Section 4.1), or using a classifier trained in a supervised manner (see Appendix E). Alternatively, we can use an *unsupervised* approach to discover semantic groups (i.e. the support of Y_t) and estimate the distribution of $Y_t \mid A_t, S_t$. For unsupervised abstractions, we use an auxiliary LLM for automated concept discovery and classification, as described in Appendix F.




Figure 2: SCM of an LM agent with abstraction variable Y

3.1 Inference method

Given a factual state s , let $a \sim A \mid s$ be the observed action (from now, we omit the time-step indices for brevity). The goal of ACF is to compute a counterfactual action A' in a different state s' given a , but without performing abduction on the token-level mechanism f_A (which is semantics-agnostic). To do so, ACF derives a counterfactual Y' for the observed abstraction value y by performing abduction over the combined mechanism $f_Y \circ f_A$. Note that ACF’s abduction step is conditioned only on the abstraction y , not the action a . Such obtained Y' represents the abstraction of the, yet unknown, counterfactual action A' . The latter is found by mapping back Y' into the action space, i.e., by deriving the posterior distribution of A given Y' and s' . In summary, ACF’s inference procedure consists of the following three steps:

1. Abduction: derive the posterior distribution of the exogenous noise for Y , $U_Y' = U_Y \mid s, y$, given the observation s, y . Formally, this is given by

$$P_{U_Y'}(u_Y' \mid s, y) = \frac{P(y \mid s, u_Y') \cdot P_{U_Y}(u_Y')}{P(y \mid s)}$$

where $P(y \mid s, u_Y')$ is the probability, induced by U_A (the randomness in action sampling), that s and u_Y' result in the observation y :

$$P(y \mid s, u_Y') = P(y = f_Y(f_A(s, U_A), s, u_Y'))$$

and $P(y \mid s) = \mathbb{E}_{U_Y}[P(y \mid s, U_Y)]$.

2. Counterfactual inference of Y' : For a given counterfactual state s' , we plug in the above posterior U_Y' to obtain a distribution of the counterfactual abstraction $Y' = Y \mid s', U_Y'$ as follows:

$$P_{Y'}(y' \mid s') = P(y' = f_Y(f_A(s', U_A), s', U_Y'))$$

3. Mapping Y' back into the action space: in the final step, we derive the counterfactual action A' in a way that its distribution is consistent with the distribution of the counterfactual abstraction Y' derived in step 2. First, we compute the posterior

$$P_{A'}(a' | y', s') = \frac{P_Y(y' | s', a') \cdot P_A(a' | s')}{P_Y(y' | s')}, \quad (3)$$

where $P_Y(y' | s', a')$ “weighs” the probability of action a' by the probability that a' leads to y' , with

$$\begin{aligned} P_A(a' | s') &= P(a' = f_A(s', U_A)) \\ P_Y(y' | s', a') &= P(y' = f_Y(a', s', U_Y)) \end{aligned}$$

and $P_Y(y' | s') = \sum_a P_Y(y' | s', a) \cdot P_A(a | s')$. Finally, we obtain the desired distribution of A' by marginalizing (3) over Y' :

$$P_{A'}(a' | s') = \mathbb{E}_{Y'}[P_{A'}(a' | Y', s')] \quad (4)$$

We stress that this approach allows us to perform counterfactual inference on A but without performing inference on U_A , i.e., by-passing the (token-level) mechanism f_A during the abduction step. This allows us to treat the LLM as a black box from which we only need to take samples.

In the above three steps, we normally estimate $f_A(\cdot, U_A)$ empirically by autoregressive sampling of the language model. However, if A consists of only one token (as in choice-based environments), we can use the precise softmax probabilities computed by the model.

Interventional consistency. Any well-defined counterfactual inference method should be consistent from an interventional viewpoint: a counterfactual represents the “individual-level” outcome of an intervention, and so, averaging the counterfactual distributions for each individual should yield the interventional distribution, i.e., the “population-level” outcome. We prove that our ACFs enjoy this consistency property, as stated below.

Proposition 1 *For some action a' , let $P_A(a' | s')$ be its interventional distribution and $P_{A'}(a' | s')$ be its ACF distribution given we observed y and s , i.e., obtained by using the posterior $U'_Y = U_Y | s, y$. Then, it holds that*

$$\mathbb{E}_{U'_Y \sim P_{U_Y}}[P_{A'}(a' | s')] = P_A(a' | s'),$$

where the expectation is taken over P_{U_Y} , the prior distribution of U_Y .

The proof is provided in Appendix A.

3.2 Interpretation and desiderata of the abstraction distribution

As discussed, we can think of the abstraction distribution as a context-dependent proxy which allows us to generate more meaningful counterfactuals. Additionally, the abstraction distribution can be seen as a (soft) partition of the outcomes into groups (e.g., cautious vs courageous actions in Figure 1). In this sense, ACF infers group-conditional outcomes rather than individual ones.

While our method is compatible with arbitrary distributions for Y , it is important that they satisfy certain intuitive criteria for meaningful counterfactuals. First, we want some degree of statistical dependency: $Y \not\perp (s, a)$, for Y to meaningfully characterise the action and the context. Second, the abstraction should not be too coarse or too fine-grained: with a small number of abstraction classes, we may oversimplify and omit important individual nuances; with too many classes, we risk introducing unnecessary complexity and reducing interpretability. In the case of supervised abstraction, achieving this balance relies on expert judgment, whereas in the unsupervised case, it depends on the concept discovery method employed.

4 Evaluation

In this section, we evaluate ACF against the token-level counterfactual inference approach of [27]. The primary goal of this evaluation is to assess how well ACF maintains high-level semantic consistency across factual and counterfactual scenarios. Our evaluation is conducted on three datasets: the MACHIAVELLI benchmark [23], the Bios dataset [8], and the GoEmotions dataset [9]. Due to

the deterministic nature of the MACHIAVELLI environment, we only present illustrative cases demonstrating our method’s effectiveness.

In text-generation settings (Sections 4.2 and 4.3), we evaluate the following metrics (explained in detail in appendix B).

1. **Abstraction Change Rate (ACR):** The proportion of instances where the most probable counterfactual abstraction value differs from the observed one. A low rate indicates that the semantic content between factual and counterfactual generations remains consistent.
2. **Counterfactual Probability Increase Rate (CPIR):** The proportion of instances where the counterfactual probability for the observed abstraction value exceeds its interventional probability. This metric evaluates whether observing a particular abstraction value increases its counterfactual probability, as desired.
3. **Semantic Tightness (ST):** The semantic similarity among different counterfactual samples generated from the same factual setting (measured via the cosine similarity of their embeddings). High semantic tightness indicates that counterfactual samples remain similar. To compare the ST values of ACF and TLCHF, we report the proportion of times ACF has better ST than TLCHF (the win rate) and the t-statistic of a paired T test.

4.1 MACHIAVELLI

Our first case study focuses on the ‘MACHIAVELLI’ benchmark [23]. This consists of a collection of text-based ‘Choose-Your-Own-Adventure’ games, extensively annotated with behavioural tendencies displayed in each scenario. In particular, these annotations measure the agents’ tendencies towards unethical (Machiavellian) behaviour. Game scenarios (states) are presented as strings of text. Available actions are defined in each scenario, and are presented to the agent as multiple-choice selections. Our agent is implemented by the OLMo-1B LLM [12]. Given a scenario and a compatible action, the transition function is deterministic, so we can evaluate the annotations of this state-action pair by looking at those of the induced next state. We use a subset of the annotations available as ‘abstractions’ for our method, characterising the agent’s actions in terms of its tendency towards physical harm, dishonesty or power seeking, to name a few (for the full list, see Appendix G). These annotations are fixed, making $P_Y(Y \mid s, a)$ a degenerate distribution.

Figure 3 compares abstract vs. token-level counterfactuals on an extracted scene from Machiavelli [23], showing how the abstract counterfactual derived distribution is more consistent with the observed annotation. (More such examples are included in Appendix H.) As we can see, performing counterfactual inference at the token level ignores the different meanings of the presented options, and instead focuses on the action labels presented (i.e., the tokens ‘0’, ‘1’ ... which are mapped to the options). This complicates the counterfactual inference, especially in scenarios where factual and counterfactual scenarios might present us different action spaces, with no exact correspondence between options. In addition, token level methods are not well defined when the cardinality of the presented action space varies. This is because the Gumbel noise associated with options that are not present in the factual setting is undefined. For our comparison, we pad the Gumbel noise vector with 0 values when the counterfactual action space is larger than the factual one, and we truncate them when the opposite is true.

4.2 Latent space interventions - gender steering

Following Ravfogel et al. [27], we investigate gender steering interventions in GPT-2-XL by modifying its latent representations. Using the MiMiC intervention [31], we learn a linear transformation that aligns the mean and covariance of male-focused biographies (source distribution) with those of female-focused biographies (target distribution), training on the gender-annotated Bios dataset [8]. While Ravfogel et al. [27] report that such transformations can unintentionally change other attributes—most notably protagonists’ professions—due to biases in both the language model and training data, our ACF method addresses these side effects by conditioning on high-level semantic concepts and explicitly maintaining consistency in the targeted abstraction.

In this setting, an intervention consists in replacing the factual state s with $s' = (\mathbf{x}, \theta')$ where the model parameters have been modified according to the MiMiC [31] transformation and the prompt has been left unchanged. This prompt $\mathbf{x}$ corresponds to the first 8 tokens of the biography the model




Figure 3: MACHIAVELLI case study. Action distributions for factual (left) and counterfactual (right) settings. The latter are obtained via the *abstract counterfactual* method, and the token level method. The observed abstraction value is ‘physical harm: 1’, associated with action ‘0’ in the factual setting. Our method correctly increases the counterfactual probabilities associated with actions which also lead to ‘physical harm: 1’. Token level counterfactual inference simply increases the probability associated with the observed action token ‘0’, without considering its high-level meaning. The Gumbel noise for token ‘3’, not present in the action factual action space, is undefined.

is to complete, following [27]. As described above, an action $a \sim A$ represents the full continuation of the biography generated by sampling from the LM until an end-of-sequence token or a fixed maximum length. We run our method both with an abstraction learned in an unsupervised manner (described in appendix F) as well as a supervised one. We define the supervised abstraction as a categorical distribution over the set of professions available (taken from the Bios dataset [8]), and we train a classifier $P_Y(Y | a, s)$ on these labels (more details on supervised abstractions in appendix E). Using this learned abstraction in the procedure defined above results in counterfactual generated texts that may shift in aspects such as phrasing or stylistic details, but ensure consistency in their higher-level semantic content. We evaluate our method on a random sample of 250 biographies, and observe in table 1 that ACF exhibits much higher abstraction value consistency from factual to counterfactual settings both with supervised and unsupervised abstractions compare to the token-level alternative.

Our method can be seen as ‘steering’ text generation towards samples that yield a specific abstraction distribution. A potential concern is that it may reverse the intervention’s effects, steering the generation back to a male-focused setting to match the factual (male-focused) distribution. However, this occurs rarely in practice. We compare the intervention’s effectiveness with the method from Ravfogel et al. [27], demonstrating similar gendered pronoun distributions across both methods over the same sample (Figure 5).

4.3 Token level interventions - emotion tracking

In the context of counterfactual text generation, one of the simplest interventions we can perform is replacing one or more tokens in the prompt $\mathbf{x} \leftarrow \mathbf{x}'$. We want to obtain a counterfactual continuation for some alternative prompt $\mathbf{x}'$ (perhaps differing only in a specific token), having observed the factual continuation a and its abstraction value y , for the factual prompt $\mathbf{x}$. Concretely, we replace the last (non-padding) token of the provided prompt with the most likely token (predicted by the model) excluding the one present in the prompt. A fitting case study for this setting, is the generation of text which conveys a specific sentiment or emotion which we can capture through Y . For this, we use the GoEmotions [9] dataset. This consists of a manually annotated dataset of 58k English

Table 1: Metrics for latent space interventions on Bios dataset, comparing *abstract counterfactuals* (ACF) with *token-level counterfactuals* (TLCF). For the ST metric, all reported values have $p < 0.001$. The paired t -statistic for the unsupervised abstraction is $t = 13.04$; the supervised abstraction is $t = 10.95$. In the ST row, we report the win rate of ACF over TLCF.

Metric	Supervised Abstraction (profession)		Unsupervised Abstraction	
	GPT2-XL		GPT2-XL	
	ACF	TLCF	ACF	TLCF
ACR ↓	0.04	0.40	0.12	0.38
CPIR ↑	0.98	0.59	0.98	0.73
ST ↑	0.78		0.81	

Reddit comments, each labelled with one or more of 27 emotion categories or ‘neutral’. In line with the framework defined above, we learn an abstraction distribution P_Y that captures the distribution of emotions in the generated text. This is implemented as a classifier trained over the GoEmotions dataset, which gives us probabilities over the 28 categories available. More details on the architecture of this classifier is available in appendix E. We evaluate our method on the GPT2-XL [25] and Llama3.2-1B [11] LLMs. Similarly to the latent space intervention study 4.2, we measure the rate of abstraction value change as described in appendix B. Table 2 show that our method achieves higher abstraction consistency than the token level one adapted from Ravfogel et al. [27], significantly decreasing the rate of change in abstraction values. We also perform the same pairwise comparison of *semantic tightness* as in the previous section, observing in all cases higher scores for ACF.

Table 2: Metrics for token level interventions on GoEmotions dataset comparing *abstract counterfactuals* (ACF) with *token-level counterfactuals* (TLCF). For the ST metric, all reported values have $p < 0.001$. With GPT2-XL, the paired t -statistic for the unsupervised abstraction variant is $t = 13.6$, and the supervised variant is $t = 10.16$. For Llama-3.2-1B, the unsupervised variant gives $t = 11.8$, and the supervised method gives $t = 8.4$. In the ST row, we report the win rate of ACF over TLCF.

Metric	Supervised Abstraction (emotion)				Unsupervised Abstraction			
	GPT2-XL		Llama-3.2-1B		GPT2-XL		Llama-3.2-1B	
	ACF	TLCF	ACF	TLCF	ACF	TLCF	ACF	TLCF
ACR ↓	0.02	0.32	0.05	0.37	0.27	0.54	0.41	0.67
CPIR ↑	0.96	0.68	0.97	0.67	0.87	0.48	0.75	0.47
ST ↑	0.76		0.72		0.82		0.80	

5 Conclusion

In this paper, we introduced *abstract counterfactuals*, a novel framework tailored for generating meaningful counterfactuals for language model agents. Our approach overcomes limitations of token-level methods by leveraging high-level semantic abstractions that capture user-relevant features. By reasoning through abstracted concepts rather than individual tokens, our method ensures consistent and interpretable counterfactual reasoning across varying contexts. Experimental evaluations on text-based games and text-generation tasks with latent-space and prompt-level interventions demonstrated the effectiveness of ACF in a wide range of settings.

Limitations As our method requires sampling from a black-box language model, the most important limitation is the computational cost of taking several samples. In addition to that, another limitation

stems from the necessity of defining an abstraction distribution, which might be tricky for certain settings. We mitigate this limitation by introducing unsupervised abstractions (see appendix F).

6 Funding Acknowledgement

Authors of this work were supported by the Engineering and Physical Sciences Research Council grants numbers EP/Y003187/1, UKRI849 and MCPS-VeriSec EP/W014785/2.

References

- [1] Kenza Amara, Rita Sevastjanova, and Mennatallah El-Assady. Challenges and Opportunities in Text Generation Explainability, May 2024. URL <http://arxiv.org/abs/2405.08468>. arXiv:2405.08468 [cs].
- [2] Onur Atan, William R. Zame, Qiaojun Feng, and Mihaela van der Schaar. Constructing Effective Personalized Policies Using Counterfactual Inference from Biased Data Sets with Many Features, July 2018. URL <http://arxiv.org/abs/1612.08082>. arXiv:1612.08082 [stat].
- [3] Leo Breiman. Statistical Modeling: The Two Cultures (with comments and a rejoinder by the author). *Statistical Science*, 16(3):199–231, 08 2001. doi: 10.1214/ss/1009213726. URL <http://dx.doi.org/10.1214/ss/1009213726>.
- [4] Anthony Brohan, Noah Brown, Justice Carbajal, Yevgen Chebotar, Xi Chen, Krzysztof Choromanski, Tianli Ding, Danny Driess, Avinava Dubey, Chelsea Finn, Pete Florence, Chuyuan Fu, Montse Gonzalez Arenas, Keerthana Gopalakrishnan, Kehang Han, Karol Hausman, Alexander Herzog, Jasmine Hsu, Brian Ichter, Alex Irpan, Nikhil Joshi, Ryan Julian, Dmitry Kalashnikov, Yuheng Kuang, Isabel Leal, Lisa Lee, Tsang-Wei Edward Lee, Sergey Levine, Yao Lu, Henryk Michalewski, Igor Mordatch, Karl Pertsch, Kanishka Rao, Krista Reymann, Michael Ryoo, Grecia Salazar, Pannag Sanketi, Pierre Sermanet, Jaspier Singh, Anikait Singh, Radu Soricut, Huang Tran, Vincent Vanhoucke, Quan Vuong, Ayzaan Wahid, Stefan Welker, Paul Wohlhart, Jialin Wu, Fei Xia, Ted Xiao, Peng Xu, Sichun Xu, Tianhe Yu, and Brianna Zitkovich. RT-2: Vision-Language-Action Models Transfer Web Knowledge to Robotic Control, July 2023. URL <http://arxiv.org/abs/2307.15818>. arXiv:2307.15818 [cs].
- [5] Lars Buesing, Theophane Weber, Yori Zwols, Sebastien Racaniere, Arthur Guez, Jean-Baptiste Lespiau, and Nicolas Heess. Woulda, Coulda, Shoulda: Counterfactually-Guided Policy Search. *arXiv preprint arXiv:1811.06272*, 2018.
- [6] Ruth M. J. Byrne. Counterfactuals in Explainable Artificial Intelligence (XAI): Evidence from Human Reasoning. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence*, pages 6276–6282, Macao, China, August 2019. International Joint Conferences on Artificial Intelligence Organization. ISBN 978-0-9992411-4-1. doi: 10.24963/ijcai.2019/876. URL <https://www.ijcai.org/proceedings/2019/876>.
- [7] Ivi Chatzi, Nina Corvelo Benz, Eleni Straitouri, Stratis Tsirtsis, and Manuel Gomez-Rodriguez. Counterfactual Token Generation in Large Language Models, September 2024. URL <http://arxiv.org/abs/2409.17027>. arXiv:2409.17027 [cs].
- [8] Maria De-Arteaga, Alexey Romanov, Hanna Wallach, Jennifer Chayes, Christian Borgs, Alexandra Chouldechova, Sahin Geyik, Krishnaram Kenthapadi, and Adam Tauman Kalai. Bias in Bios: A Case Study of Semantic Representation Bias in a High-Stakes Setting, January 2019. URL <http://arxiv.org/abs/1901.09451>. arXiv:1901.09451.
- [9] Dorottya Demszky, Dana Movshovitz-Attias, Jeongwoo Ko, Alan Cowen, Gaurav Nemade, and Sujith Ravi. GoEmotions: A Dataset of Fine-Grained Emotions, June 2020. URL <http://arxiv.org/abs/2005.00547>. arXiv:2005.00547 [cs].
- [10] Sebastian Farquhar, Jannik Kossen, Lorenz Kuhn, and Yarin Gal. Detecting Hallucinations in Large Language Models using Semantic Entropy. *Nature*, 630(8017):625–630, June 2024. ISSN 1476-4687. doi: 10.1038/s41586-024-07421-0. URL <https://www.nature.com/articles/s41586-024-07421-0>. Publisher: Nature Publishing Group.

- [11] Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, et al. The llama 3 Herd of Models. *arXiv preprint arXiv:2407.21783*, 2024.
- [12] Dirk Groeneveld, Iz Beltagy, Pete Walsh, Akshita Bhagia, Rodney Kinney, Oyvind Tafjord, Ananya Harsh Jha, Hamish Ivison, Ian Magnusson, Yizhong Wang, Shane Arora, David Atkinson, Russell Authur, Khyathi Raghavi Chandu, Arman Cohan, Jennifer Dumas, Yanai Elazar, Yuling Gu, Jack Hessel, Tushar Khot, William Merrill, Jacob Morrison, Niklas Muennighoff, Aakanksha Naik, Crystal Nam, Matthew E. Peters, Valentina Pyatkin, Abhilasha Ravichander, Dustin Schwenk, Saurabh Shah, Will Smith, Emma Strubell, Nishant Subramani, Mitchell Wortsman, Pradeep Dasigi, Nathan Lambert, Kyle Richardson, Luke Zettlemoyer, Jesse Dodge, Kyle Lo, Luca Soldaini, Noah A. Smith, and Hannaneh Hajishirzi. Olmo: Accelerating the Science of Language Models, 2024. URL <https://arxiv.org/abs/2402.00838>.
- [13] Taicheng Guo, Xiuying Chen, Yaqi Wang, Ruidi Chang, Shichao Pei, Nitesh V. Chawla, Olaf Wiest, and Xiangliang Zhang. Large Language Model based Multi-Agents: A Survey of Progress and Challenges, April 2024. URL <http://arxiv.org/abs/2402.01680>. arXiv:2402.01680 [cs].
- [14] Joseph Y. Halpern. Cause, Responsibility, and Blame: A Structural-Model Approach, December 2014. URL <http://arxiv.org/abs/1412.2985>. arXiv:1412.2985 [cs].
- [15] Joseph Y Halpern and Judea Pearl. Causes and Explanations: A Structural-Model Approach. Part II: Explanations. *Br. J. Philos. Sci.*, 56(4):889–911, 2005.
- [16] Milad Kazemi, Jessica Lally, Ekaterina Tishchenko, Hana Chockler, and Nicola Paoletti. Counterfactual Influence in Markov Decision Processes, February 2024. URL <http://arxiv.org/abs/2402.08514>. arXiv:2402.08514.
- [17] Jinqi Lai, Wensheng Gan, Jiayang Wu, Zhenlian Qi, and Philip S. Yu. Large Language Models in Law: A Survey. 2023. doi: 10.48550/ARXIV.2312.03718. URL <https://arxiv.org/abs/2312.03718>. Publisher: arXiv Version Number: 1.
- [18] Chris J Maddison, Daniel Tarlow, and Tom Minka. A* sampling. *Advances in neural information processing systems*, 27, 2014.
- [19] Prashan Madumal, Tim Miller, Liz Sonenberg, and Frank Vetere. Explainable Reinforcement Learning Through a Causal Lens, November 2019. URL <http://arxiv.org/abs/1905.10958>. arXiv:1905.10958 [cs].
- [20] Nikita Mehandru, Brenda Y. Miao, Eduardo Rodriguez Almaraz, Madhumita Sushil, Atul J. Butte, and Ahmed Alaa. Evaluating Large Language Models as Agents in the Clinic. *npj Digital Medicine*, 7(1):84, April 2024. ISSN 2398-6352. doi: 10.1038/s41746-024-01083-y. URL <https://www.nature.com/articles/s41746-024-01083-y>.
- [21] Roberto Navigli, Simone Conia, and Björn Ross. Biases in Large Language Models: Origins, Inventory, and Discussion. *J. Data and Information Quality*, 15(2):10:1–10:21, June 2023. ISSN 1936-1955. doi: 10.1145/3597307. URL <https://dl.acm.org/doi/10.1145/3597307>.
- [22] Michael Oberst and David Sontag. Counterfactual Off-Policy Evaluation with Gumbel-Max Structural Causal Models. In *International Conference on Machine Learning*, pages 4881–4890. PMLR, 2019.
- [23] Alexander Pan, Jun Shern Chan, Andy Zou, Nathaniel Li, Steven Basart, Thomas Woodside, Jonathan Ng, Hanlin Zhang, Scott Emmons, and Dan Hendrycks. Do the Rewards Justify the Means? Measuring Trade-Offs Between Rewards and Ethical Behavior in the Machiavelli Benchmark. *ICML*, 2023.
- [24] Judea Pearl. *Causality*. Cambridge University Press, New York, 2009.
- [25] Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. Language Models are Unsupervised Multitask Learners. *OpenAI blog*, 1(8):9, 2019.

- [26] Shauli Ravfogel, Anej Svete, Vésteinn Snæbjarnarson, and Ryan Cotterell. Counterfactual Generation from Language Models, November 2024. URL <http://arxiv.org/abs/2411.07180>. arXiv:2411.07180.
- [27] Shauli Ravfogel, Anej Svete, Vésteinn Snæbjarnarson, and Ryan Cotterell. Gumbel Counterfactual Generation From Language Models. In *The Thirteenth International Conference on Learning Representations*, 2025. URL <https://openreview.net/forum?id=TUC0ZT2zIQ>.
- [28] Nils Reimers and Iryna Gurevych. Sentence-BERT: Sentence Embeddings using Siamese BERT-Networks. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, 11 2019. URL <https://arxiv.org/abs/1908.10084>.
- [29] Juan-Pablo Rivera, Gabriel Mukobi, Anka Reuel, Max Lamparth, Chandler Smith, and Jacquelyn Schneider. Escalation Risks from Language Models in Military and Diplomatic Decision-Making. In *The 2024 ACM Conference on Fairness, Accountability, and Transparency*, pages 836–898, Rio de Janeiro Brazil, June 2024. ACM. ISBN 9798400704505. doi: 10.1145/3630106.3658942. URL <https://dl.acm.org/doi/10.1145/3630106.3658942>.
- [30] Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. DistilBERT, a Distilled Version of BERT: Smaller, Faster, Cheaper and Lighter, 2020. URL <https://arxiv.org/abs/1910.01108>.
- [31] Shashwat Singh, Shauli Ravfogel, Jonathan Herzig, Roei Aharoni, Ryan Cotterell, and Ponurangam Kumaraguru. Representation Surgery: Theory and Practice of Affine Steering, July 2024. URL <http://arxiv.org/abs/2402.09631>. arXiv:2402.09631 [cs].
- [32] Zhongxiang Sun. A Short Survey of Viewing Large Language Models in Legal Aspect. 2023. doi: 10.48550/ARXIV.2303.09136. URL <https://arxiv.org/abs/2303.09136>. Publisher: arXiv Version Number: 1.
- [33] Xiangru Tang, Anni Zou, Zhuosheng Zhang, Ziming Li, Yilun Zhao, Xingyao Zhang, Arman Cohan, and Mark Gerstein. MedAgents: Large Language Models as Collaborators for Zero-shot Medical Reasoning, June 2024. URL <http://arxiv.org/abs/2311.10537>. arXiv:2311.10537 [cs].
- [34] Jin Tian and Judea Pearl. Probabilities of Causation: Bounds and Identification. *Ann. Math. Artif. Intell.*, 28(1-4):287–313, 2000.
- [35] Alexander Sasha Vezhnevets, John P Agapiou, Avia Aharon, Ron Ziv, Jayd Matyas, Edgar A Duéñez-Guzmán, William A Cunningham, Simon Osindero, Danny Karmon, and Joel Z Leibo. Generative Agent-Based Modeling with Actions Grounded in Physical, Social, or Digital Space using Concordia. *arXiv preprint arXiv:2312.03664*, 2023.
- [36] Guanzhi Wang, Yuqi Xie, Yunfan Jiang, Ajay Mandlekar, Chaowei Xiao, Yuke Zhu, Linxi Fan, and Anima Anandkumar. Voyager: An Open-Ended Embodied Agent with Large Language Models, October 2023. URL <http://arxiv.org/abs/2305.16291>. arXiv:2305.16291 [cs].
- [37] Lei Wang, Chen Ma, Xueyang Feng, Zeyu Zhang, Hao Yang, Jingsen Zhang, Zhiyuan Chen, Jiakai Tang, Xu Chen, Yankai Lin, Wayne Xin Zhao, Zhewei Wei, and Ji-Rong Wen. A Survey on Large Language Model based Autonomous Agents. *Frontiers of Computer Science*, 18(6): 186345, December 2024. ISSN 2095-2228, 2095-2236. doi: 10.1007/s11704-024-40231-1. URL <http://arxiv.org/abs/2308.11432>. arXiv:2308.11432 [cs].
- [38] Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik Narasimhan, and Yuan Cao. React: Synergizing Reasoning and Acting in Language Models. In *International Conference on Learning Representations (ICLR)*, 2023.
- [39] Xizhou Zhu, Yuntao Chen, Hao Tian, Chenxin Tao, Weijie Su, Chenyu Yang, Gao Huang, Bin Li, Lewei Lu, Xiaogang Wang, Yu Qiao, Zhaoxiang Zhang, and Jifeng Dai. Ghost in the Minecraft: Generally Capable Agents for Open-World Environments via Large Language Models with Text-based Knowledge and Memory, June 2023. URL <http://arxiv.org/abs/2305.17144>. arXiv:2305.17144 [cs].

A Proof of Proposition 1

We want to show that the expectation of our counterfactual distribution w.r.t. the exogenous values distribution is equal to the interventional distribution. Formally,

$$\mathbb{E}_{U'_{Y'} \sim P_{U_Y}} [P_{A'}(a' | s')] = P_A(a' | s')$$

For simplicity of notation, we omit conditioning on s' . Also, for simplicity of notation, we assume all variables are discrete. The proof below would work by replacing sums with integrals for continuous variables. So, the LHS of the above statement can be rewritten as

$$\sum_{u'} \sum_{y'} P(u') \cdot P(y' | u') \frac{P(y' | a') \cdot P(a')}{P(y')}$$

and we want to show this is equal to $P(a')$. By noting that (by Bayes rule) $\frac{P(y' | u')}{P(y')} = \frac{P(u' | y')}{P(u')}$, then the above expression simplifies to

$$\sum_{u'} \sum_{y'} P(u' | y') \cdot P(y' | a') \cdot P(a')$$

which is equivalent to

$$\sum_{y'} P(y' | a') \cdot P(a') \cdot \sum_{u'} P(u' | y')$$

By Bayes rule, we can rewrite $\sum_{u'} P(u' | y')$ as $\frac{1}{P(y')} \sum_{u'} P(u', y')$ and it's easy to see that this equals 1. So, the above expression simplifies to

$$\sum_{y'} P(y' | a') \cdot P(a') = \sum_{y'} P(y', a')$$

which is equal to $P(a')$.

B Metrics

1. **Abstraction Change Rate:** This metric calculates the proportion of instances where the most probable counterfactual abstraction value differs from the observed abstraction value. Formally, it calculates the number of instances where:

$$\operatorname{argmax}_{y' \in \mathcal{Y}'} \left\{ \frac{1}{|\mathbf{a}'|} \sum_{a \in \mathbf{a}'} P_{cf}(y' | a, s') \right\} \neq y$$

Where y is the abstraction value observation. A low abstraction change rate indicates that the semantic content of the counterfactual generation remains consistent with the initial generation.

2. **Counterfactual Probability Increase Rate:** This metric measures the proportion of cases where the counterfactual probability for the observed abstraction value is greater than its interventional probability. Formally, it calculates the number of instances where:

$$P_{cf}(y | a', s') > P(y | a, s')$$

This metric evaluates whether the counterfactual probability derived from counterfactual samples ($a' \sim P_{A'}(a' | s')$) exceeds that of interventional samples ($a \sim P_A(a | s')$).


3. **Semantic Tightness:** We also evaluate the semantic similarity between different counterfactual samples generated from the same factual setting. Formally, given a set of strings $\mathbf{a} = \{a_1, a_2, \dots, a_n\}$ and a semantic embedding model λ , we can compute the semantic tightness as:

$$\text{semantic_tightness}(\mathbf{a}) = \frac{1}{|\mathbf{a}|^2} \sum_{i=1}^{|\mathbf{a}|} \sum_{j=1}^{|\mathbf{a}|} \cos_sim(\lambda(a_i), \lambda(a_j))$$

where we measure the average cosine similarity between all pairs of embeddings from the strings in $\mathbf{a}$. As our embedding model λ , we use the 'all-mpnet-base-v2' model from the 'sentence-transformers' library [28]. High semantic tightness indicates that counterfactual samples remain similar, even when generated independently from the same factual context.


C Counterfactual sample semantic tightness

Figure 4: Scatterplots showing the semantic tightness of counterfactual samples generated with *abstract counterfactuals* and *token-level counterfactuals* for the same initial state s . Each point represents an initial state.



D Gendered pronoun distributions

Figure 5: Pronouns distribution over counterfactual samples generated with our method and token level method [27] for the gender steering intervention on GPT2-XL



E Supervised LLM abstraction

For both text-generation tasks we run experiments with supervised abstractions. Here, a classifier models the distribution $P_Y(Y | s, a)$, after being trained on an annotated dataset. This classifier is implemented by fine-tuning a DistilBERT [30] language model on the respective dataset. For the ‘gender steering’ latent space interventions we fine-tune the model to predict the protagonist’s profession, using the BiosBias [8] dataset, resulting in a model with an f1 score of 0.85. The available profession categories from this dataset are: *professor, physician, attorney, photographer, journalist, nurse, psychologist, teacher, dentist, surgeon, architect, painter, model, poet, filmmaker, software engineer, accountant, composer, dietitian, comedian, chiropractor, pastor, paralegal, yoga teacher, dj, interior designer, personal trainer, rapper*.

In the case of the token replacement interventions, we fine-tune the classifier on the GoEmotions dataset [9], assigning the generated text to one of the following categories: *admiration, amusement, anger, annoyance, approval, caring, confusion, curiosity, desire, disappointment, disapproval, disgust, embarrassment, excitement, fear, gratitude, grief, joy, love, nervousness, neutral, optimism, pride, realization, relief, remorse, sadness, surprise, neutral*. The emotion classification pipeline has an f1 score of 0.63.

F Unsupervised LLM abstractions

Our goal is to discover an effective sample space for Y which captures the meaning of generated text completions. This is similar to the semantic clustering performed by [10] to compute entropy at the semantic level. For semantic grouping, however, we do not use the bi-directional entailment algorithm proposed by [10], as it is more suited for specific claims (i.e. Q&A format) rather than prompted text generation. Instead, we use a two-stage process, which consists in a *topic discovery* phase, and a *topic assignment*. Given a factual and interventional states s, s' , we sample text continuations from both, and join them in a single set $\mathbf{a}$. This set is fed to a language model prompted to identify a list of topics that partition the set. Each sample $a \in \mathbf{a}$ is then fed individually to a language model prompted to assign it the most likely topic from the previously discovered set. We use the model’s logits over topic indices to construct the distributions of $Y | a$.

For the experiments throughout the paper we use ‘gpt-4o-mini’ as language model. As the openai API (as of this writing) only provides logits the 20 most likely tokens, we filter out invalid token id responses and add an ‘Other’ category to the sample space of Y after *topic discovery*.

Unsupervised LLM Abstraction Prompts

Topic Discovery Prompt:

"You are a helpful assistant, and you output python code
Your task is to come up with groups that partition the following list of strings at the semantic level.
<formatted string list>
You must answer exclusively with a valid list of tuples, where element 0 is a category title (string), and element 1 is a very brief description of that category"

Topic Assignment Prompt:

"You are a helpful assistant, and you only output integers within the range specified by the number of topics.
Given a set of topics, their descriptions, and their IDs (0-indexed) and a new string, assign the string to the most appropriate topic. Reply only with the ID of the topics.
Topics: <formatted topics>
String to be classified: <formatted string>
topic ID:"

G MACHIAVELLI annotations

We are interested in the annotations for the ‘morality’ category from the MACHIAVELLI [23] benchmark. The available (binary) annotations, for each scene, are the following: *killing, physical harm, non physical harm, intending harm, deception, manipulation, promises, stealing, trespassing, spying, vandalism, fairness, and other*.

H MACHIAVELLI further examples

Figures 6, 7, 8 show further examples of *abstract counterfactuals* in MACHIAVELLI [23] settings with ‘morality’ annotations as abstractions.




Figure 6: Scenes from the ‘Aegis Project Files 180’ game, abstraction ‘manipulation’.




Figure 7: Scenes from the ‘Aegis Project Files 180’ game, abstraction ‘physical harm’.




Figure 8: Scenes from the ‘Alexandria’ game, abstraction ‘killing’.

I Gender steering examples

Gender Steering Example

Factual Text:

Also a traditional reporter, this New Jersey native went on to do more than 40 stories for the News of the World, then founded the news website MuckRock to crowdsource information and ultimately went on to co-found the tech-oriented think tank, the Future of Privacy Forum. Now he's helping others get their stories out by sharing his own. He's a full member of the PRweb team, as we are proud to be [...]

Predicted Abstraction: 'Journalist'

Token-Level Counterfactual Text:

Also a traditional reporter, this New Jersey native went to Boston College and USC before earning a law degree. She is currently managing editor of the weekly column "Sex and Shared Desires," and occasionally contributes to "TribLive" on the Huffington Post, LGBT Nation!

Predicted Abstraction: 'Attorney'

Abstract Counterfactual Text:

Also a traditional reporter, this New Jersey-born reporter is among those chosen to represent KQED's public television reporting. Ms. Williams has reported from across the country, from Israel and Moscow, and on eight continents. She has been with the Public Insight Network for two decades. Before joining KQED, Ms. Williams was a senior producer at member station WITN in Washington, DC [...]

Predicted Abstraction: 'Journalist'

J Compute Resources

We run all the reported experiments on a server equipped with an x86_64, 128-core CPU with 405.2 GB of RAM and an NVIDIA A40 GPU with 48GB of VRAM. The server runs Ubuntu 20.04.6 LTS.

Counterfactual sample generation pipelines take around 4 hours each (for each configuration of model, abstraction type and dataset). Fine-tuning of the supervised abstraction models takes around 1h for each abstraction.

NeurIPS Paper Checklist

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [Yes]

Justification: The method is introduced in section 3 and experimental results are shown in section 4.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [Yes]

Justification: Limitations are discussed in section 5. One of the limitations (the choice of an abstraction distribution) is addressed and mitigated by the introduction of unsupervised abstractions (see appendix F).

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory assumptions and proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [Yes]

Justification: The only theoretical result (proof of consistency of our counterfactual method) is presented in appendix A.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental result reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: The method we introduce is explained in detail in multiple steps in section 3. Furthermore, we include practical details about the construction of the supervised and unsupervised abstractions (including LLM prompts) respectively in appendices E and F.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
 - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
 - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
 - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [Yes]

Justification: We provide a .zip file as part of the submission with the full codebase required for running and evaluating the results, as well as documentation on how to use it.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental setting/details

Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: We provide extensive details about experimental settings, configurations, models and data in the paper appendices, as well as in the provided anonymized codebase.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. Experiment statistical significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [No]

Justification: While we do report the statistical significance of the semantic tightness metric, we do not do so for the other metrics, due to limits in compute budget and API costs.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.

- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments compute resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [\[Yes\]](#)

Justification: Compute resources and time of execution for experiments are detailed in appendix J.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code of ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines?>

Answer: [\[Yes\]](#)

Justification: We have reviewed and adhered to the NeurIPS Code of Ethics.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [\[Yes\]](#)

Justification: We discuss social impacts of language model agents and specifically counterfactual inference in the introduction of our work. We do not foresee negative societal impacts arising from counterfactual inference. Our work addresses the possibility of biased data and language models resulting in biased counterfactuals, and our research aims at limiting this.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: We are not releasing any new data or models. We do not foresee any misuse risk from the released method and experiment code.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification: We cite the original owners of all assets and code we build upon, and we report their licenses in footnotes.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.

- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. **New assets**

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [Yes]

Justification: The only new asset produced as part of this paper is the associated code. This is provided in an anonymized .zip file, including documentation about its functioning and usage.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and research with human subjects**

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: Our reserach didn't include any crowdsourcing nor research with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional review board (IRB) approvals or equivalent for research with human subjects**

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: This paper does not involve crowdsourcing nor research with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. **Declaration of LLM usage**

Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does not impact the core methodology, scientific rigorousness, or originality of the research, declaration is not required.

Answer: [\[Yes\]](#)

Justification: Aside from studying LLMs themselves throughout the paper, we use LLMs as components of our method, implementing the ‘abstraction’ distributions. We describe their usage in this regard in detail in appendices E and F.

Guidelines:

- The answer NA means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
- Please refer to our LLM policy (<https://neurips.cc/Conferences/2025/LLM>) for what should or should not be described.