

LIBERTY: A Causal Framework for Benchmarking Concept-Based Explanations of LLMs with Structural Counterfactuals

Gilat Toker*, Nitay Calderon*, Ohad Amosy, Roi Reichart

Faculty of Data and Decision Sciences, Technion – Israel Institute of Technology
{gilatt, nitay}@campus.technion.ac.il, roiri@technion.ac.il

*Second author supervised the project and led the writing.

Abstract

Concept-based explanations quantify how high-level concepts (e.g., gender or experience) influence model behavior, which is crucial for decision-makers in high-stakes domains. Recent work evaluates the faithfulness of such explanations by comparing them to reference causal effects estimated from counterfactuals. In practice, existing benchmarks rely on costly human-written counterfactuals that serve as an imperfect proxy. To address this, we introduce a framework for constructing datasets containing structural counterfactual pairs: LIBERTY (LLM-based **I**nterventional **B**enchmark for **E**xplainability with **R**eference **T**argets). LIBERTY is grounded in explicitly defined Structured Causal Models (SCMs) of the text generation, interventions on a concept propagate through the SCM until an LLM generates the counterfactual. We introduce three datasets (disease detection, CV screening, and workplace violence prediction) together with a new evaluation metric, order-faithfulness. Using them, we evaluate a wide range of methods across five models and identify substantial headroom for improving concept-based explanations. LIBERTY also enables systematic analysis of model sensitivity to interventions: we find that proprietary LLMs show markedly reduced sensitivity to demographic concepts, likely due to post-training mitigation. Overall, LIBERTY provides a much-needed benchmark for developing faithful explainability methods.¹

1 Introduction

AI systems, especially Large Language Models (LLMs), increasingly drive decisions in sensitive and high-stakes domains where *textual input* plays a central role, such as finance, education, healthcare, and law (Guidotti et al., 2018; Balkir et al., 2022; Kasneci et al., 2023; Shui et al., 2023; Luo et al., 2024; Nie et al., 2024; Benkirane et al., 2024).

¹<https://github.com/GilatToker/Liberty-benchmark>

The decisions of these opaque systems are difficult to explain, making explainability a central research challenge (Guidotti et al., 2018; Balkir et al., 2022; Luo et al., 2024). Among the many approaches to explainability, *concept-based methods* are particularly relevant when the stakeholders are decision-makers and end-users (Calderon and Reichart, 2024). These methods focus on quantifying the influence of high-level, human-interpretable concepts, such as gender, race, or professional experience, on model predictions (Kim et al., 2018; Künzel et al., 2019; Yeh et al., 2020; Feder et al., 2021; Wu et al., 2022; Gat et al., 2023).

Recent studies emphasize that explanations lacking a causal basis often fail to achieve true faithfulness (Lyu et al., 2022; Gat et al., 2023; Yeo et al., 2024). In causality, a causal graph encodes concepts as variables and their relationships as edges (Pearl, 2009). This structure enables us to identify the roles of concepts, such as confounders, mediators, and colliders, and to estimate the causal effect of a target concept on the model (Abraham et al., 2022). Despite progress at the intersection of AI and causality (Wood-Doughty et al., 2018; Feder et al., 2021; Wu et al., 2023; Zhang et al., 2023), a fundamental challenge remains: evaluating whether an explanation is faithful requires comparing it to the true underlying causal mechanisms. In practice, the ground-truth mechanism is inaccessible, leaving us without a reliable benchmark for explainability methods.

One approach to address this benchmarking challenge was introduced by Abraham et al. (2022). They propose using an *interventional dataset* as a systematic framework for evaluation of explanations. In their interventional dataset, CEBaB, each test example is paired with a human-written counterfactual generated by modifying a concept. The individual causal concept effect (ICaCE) is then estimated by contrasting the model’s outputs on the original text and its counterfactual. Expla-

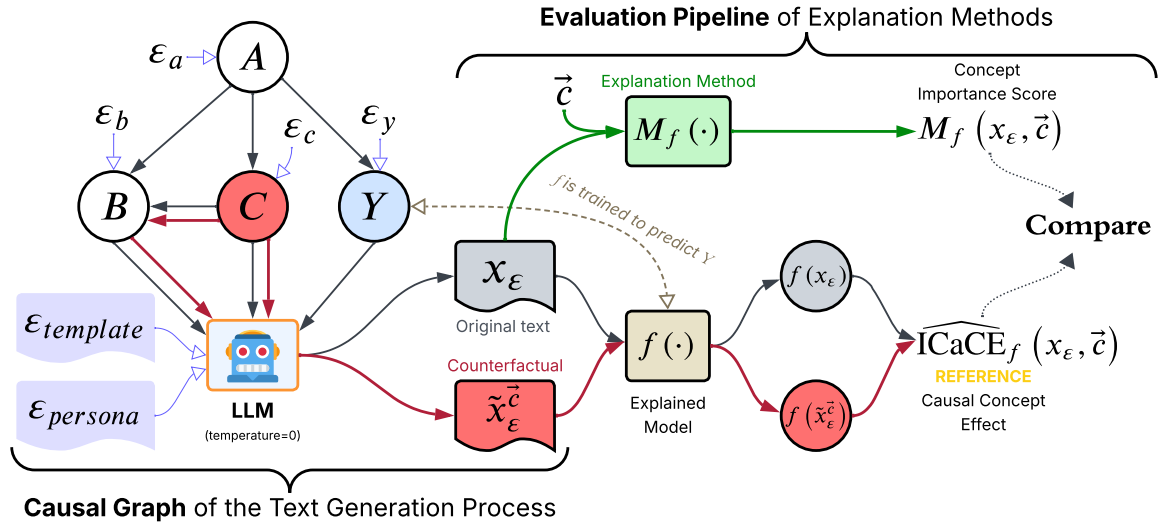


Figure 1: **Illustration of LIBERTy**: The goal is to evaluate an explanation method M_f that explains the impact of changing a concept C (by \vec{c}) on model f . **Left**: The causal graph representing the text generation process. Exogenous noise variables are denoted by ε , while the endogenous variables (in this illustration) are the concepts A, B, C, Y , the LLM-generated text x_ε , and the model prediction $f(x_\varepsilon)$. The process of generating a structural counterfactual for the change \vec{c} is highlighted in red: C is assigned a new value and propagated through the causal graph (with ε fixed) until the LLM generates the counterfactual $x_\varepsilon^{\vec{c}}$. **Right**: The explanation $M_f(x_\varepsilon, \vec{c})$ is compared against the reference individual causal concept effect (ICaCE), defined as the difference between $f(x_\varepsilon^{\vec{c}})$ and $f(x_\varepsilon)$.

nation methods are evaluated by comparing their importance score (of the concept) against the estimated causal effect. While CEBaB represents a significant step toward causal evaluation, it remains limited, especially given LLMs’ current capabilities. First, CEBaB is confined to sentiment analysis of restaurant reviews, which are short, simple texts. Second, its causal graph comprises only four concepts, with simple relationships (no hierarchical structure). Finally, the counterfactuals are written by human annotators rather than arising from actual interventions in the *data-generating process* (DGP). Consequently, the causal effect references used as “ground truth” for evaluation are themselves approximations of some unobserved effects.

In this work, we address these limitations by introducing a novel framework for generating interventional datasets with structural counterfactuals that define reference causal effects: LIBERTy (LLM-based **I**nterventional **B**enchmark for **E**xplainability with **R**eference **T**argets). LIBERTy, illustrated in Figure 1, is based on a simple yet effective idea: explicitly defining a structured causal model (SCM) for text generation. In this framework, the LLM is a component of the SCM that instantiates concepts as natural language text. To make LLM outputs more diverse and realistic, we provide it with grounding context, such as tem-

plated real-world text and author persona, which act as exogenous noise variables in the SCM. Counterfactuals are generated by intervening on a concept (assigning it a new value) and propagating this change through the SCM until the LLM produces the corresponding counterfactual. As a result, LIBERTy provides structural counterfactuals², eliminating the need for costly human annotations and ensuring alignment between the evaluation reference target and the DGP.

LIBERTy comprises three datasets, each designed around a major societal challenge: disease detection, CV screening, and workplace violence prediction. We also propose a new evaluation measure, order-faithfulness, that quantifies how well an explanation method captures the relative ordering of effects induced by concept interventions. This makes it suitable for evaluating explanation methods that provide importance scores on arbitrary scales, rather than direct causal effect estimates.

Using LIBERTy, we conduct extensive experiments to explain five NLP models and LLMs.

²Formally, LIBERTy counterfactuals are gold with respect to the DGP (which the LLM that generates text is part of): given the SCM and observed exogenous values, they are generated via Pearl’s three-step procedure (Pearl, 2013). However, since the DGP itself and the resulting texts are synthetic, we use the term silver to refer to these structural counterfactuals. Yet, as LLMs generate an increasing share of real-world data, such setups are both common and practically meaningful.

We benchmark a range of concept-based explanation methods, including linear erasure, counterfactual generation, matching, and concept attributions. Our results show that matching methods based on representations from a dataset-specific fine-tuned model perform best overall. Still, we find substantial headroom for improvement, highlighting the need for continued explainability research.

Besides evaluating explanations, LIBERTy enables us to analyze the sensitivity of each explained model to concept interventions. For example, when the model predicts a candidate’s qualification based on their CV, we examine how its prediction changes when we intervene on the candidate’s race, and compare this change to the effect specified in the SCM. Our results show that fine-tuned models can track the ground-truth effects of the data. In contrast, some LLMs (like GPT-4o) exhibit very low sensitivity to demographic concepts, potentially due to dedicated post-training alignment.

Overall, our study represents an important step toward addressing the long-standing challenge of explainability evaluation. By introducing LIBERTy, we provide researchers with a reliable, scalable, and flexible causal framework for benchmark generation, paving the way for the development of more faithful explainability methods.

2 Related Work

Unlike token-level methods (Calderon and Reichart, 2024; Luo et al., 2024; Zhao et al., 2024), concept-based explainability quantifies the influence of human-interpretable concepts, aligning better with cognitive processes (Alqaraawi et al., 2020; Kim et al., 2022; Poeta et al., 2023; Calderon and Reichart, 2024). Applications range from bias detection and healthcare to reward modeling (Cornacchia et al., 2023; Bouchacourt and Denoyer, 2019; Ghorbani et al., 2019; Calderon et al., 2025; Peled-Cohen et al., 2025). Prominent approaches include attribution (Ribeiro et al., 2016; Lundberg and Lee, 2017; Kim et al., 2018; Yeh et al., 2020), erasure (Ravfogel et al., 2022; Belrose et al., 2023), counterfactuals (Feder et al., 2021; Robeer et al., 2021; Wu et al., 2021; Gat et al., 2023), matching (Veitch et al., 2020; Zhang et al., 2023; Gat et al., 2023; Jiang et al., 2025), and bottleneck models (Koh et al., 2020; Dalvi et al., 2022; Yu et al., 2024).

Evaluating explanations is challenging due to inaccessible ground truth (Yang et al., 2019; Hedström et al., 2023; Lee et al., 2025; Seth and

Sankarapu, 2025), therefore, evaluation often relies on indirect proxies or token-level metrics (Hase and Bansal, 2020; Samek et al., 2021; Thorne et al., 2019; Wang et al., 2022; Gurrupu et al., 2023). While CEBaB (Abraham et al., 2022) introduced an interventional dataset to assess explanations, recent studies emphasize the need for benchmarks that capture the structural complexity of real-world problems and address statistical fallacies in LLMs (Chaleshtori et al., 2024; Du et al., 2025). Hopefully, LIBERTy would address these gaps.

3 Evaluation of Explanations

In this section, we provide the relevant causal background and outline our causal approach to evaluating two explanation scopes. *Local explanations* capture how a concept influences a prediction for a specific instance, whereas *global explanations* capture its influence across the entire data distribution. We evaluate explanations by comparing them with causal effects: locals against individual-level effects and globals against population-level effects.

3.1 Causality Background

Structural Causal Models (SCMs) We adopt the SCM framework of Pearl (2009). An SCM consists of exogenous and endogenous variables, together with structural equations. Each endogenous variable is defined as a function of its parent endogenous variables and its associated exogenous noise variable. The induced causal graph is a directed acyclic graph encoding these dependencies. An example is given in Figure 1, where the endogenous variables are the concepts, the LLM-generated text x_ε , and the prediction of the explained model $f(x_\varepsilon)$. The exogenous variables include Gaussian noise terms $(\varepsilon_a, \varepsilon_b, \varepsilon_c, \varepsilon_y)$, or randomly sampled auxiliary text provided to the LLM ($\varepsilon_{template}$ and $\varepsilon_{persona}$). Complementing the SCM with explicit distributions over the noise yields the *data-generating process (DGP)*.

Counterfactuals Within the SCM framework, a *counterfactual* is the outcome of an intervention that assigns a different value to a concept while keeping all the exogenous variables fixed. In our setting, a *textual counterfactual* is generated by propagating the intervened concept assignment through the SCM. A *prediction counterfactual* is obtained by passing this counterfactual text to the explained model and observing its new prediction. Because the DGP is fully specified and LLM de-

Box 2.1: Definitions: Causal Concept Effects and Estimators

Definition 1 (Causal Concept Effect (CaCE) and Individual CaCE (ICaCE)).

$$\begin{aligned}\text{CaCE}_f(\vec{c}) &= \mathbb{E}[f(X) \mid \text{do}(C = c')] - \mathbb{E}[f(X) \mid \text{do}(C = c)] \\ \text{ICaCE}_f(x_\varepsilon, \vec{c}) &= \mathbb{E}[f(X) \mid \text{do}(C = c', \mathcal{E} = \varepsilon)] - f(x_\varepsilon)\end{aligned}$$

Definition 2 (Empirical CaCE and ICaCE).

$$\begin{aligned}\widehat{\text{CaCE}}_f(\vec{c}) &= \frac{1}{|D|} \sum_{x_{\varepsilon^*} \in D} [f(\tilde{x}_{\varepsilon^*}^{c' \rightarrow c'}) - f(\tilde{x}_{\varepsilon^*}^{c^* \rightarrow c})] \\ \widehat{\text{ICaCE}}_f(x_\varepsilon, \vec{c}) &= f(\tilde{x}_\varepsilon^{\vec{c}}) - f(x_\varepsilon)\end{aligned}$$

coding is deterministic (with the temperature set to zero), the counterfactuals align with Pearl’s definition of structural counterfactuals (Pearl, 2013).

Causal Effect of Concepts We consider two levels of causal effects: the *Causal Concept Effect (CaCE)* (Goyal et al., 2019), and the *individual CaCE (ICaCE)*. Ideally, a faithful explanation method would estimate the **CaCE as a global explanation** and the **ICaCE as a local explanation** (Gat et al., 2023). Formally, let C denote the concept whose value changes from c to c' (written \vec{c}), and let \mathcal{E} be exogenous variables with ε values. Then, x_ε is the resulting text, and the prediction of the explained model is $f(x_\varepsilon)$, which is a vector with predicted softmax probabilities of each class of Y . We denote expectations under the interventional distribution by the standard do-operator notation $\mathbb{E}[\cdot \mid \text{do}(C = c')]$. The formal definitions are provided in Def 1. Both CaCE and ICaCE are vectors that capture effects on all classes.

3.2 Evaluating Explanations

Estimating Causal Effects Both the CaCE and ICaCE are theoretical quantities, and in practice, we estimate them using counterfactuals. For x_ε , we denote its counterfactual by $\tilde{x}_\varepsilon^{\vec{c}}$. The formal definitions of the estimators are provided in Def 2. In our setting, $\widehat{\text{ICaCE}}_f$ is exact because, with fixed \mathcal{E} and deterministic decoding, $\mathbb{E}[f(X) \mid \text{do}(C = c', \mathcal{E} = \varepsilon)] = f(\tilde{x}_\varepsilon^{\vec{c}})$. If decoding is stochastic (e.g., temperature > 0), additional noise is introduced through token sampling (see the discussion in Appendix A).

Evaluation Pipeline The explained model f is trained on DGP-sampled data D_f .³ The explanation method M is trained on pairs $(x, f(x)) : x \in D_M$, with optional access to gold

concept values or other auxiliary information, depending on the evaluator choice. For evaluating M_f , we use the interventional test set D_C , where C denotes the set of concept changes. For each change, $D_{\vec{c}}$ consists of pairs of textual examples and their counterfactuals, $(x_\varepsilon, \tilde{x}_\varepsilon^{\vec{c}})$. From these, we compute $\widehat{\text{CaCE}}_f(\vec{c})$ and $\{\widehat{\text{ICaCE}}_f(x, \vec{c})\}_x$, as well as the corresponding explanation scores: $M_f(\vec{c})$ for global methods and $\{M_f(x, \vec{c})\}_x$ for local ones. We next describe the evaluation measures for local explanations, which can be extended to global explanations with minor modifications.

Evaluation Measures CEBaB reports the average ICaCE Error Distance over all the concept changes, defined as the distance between the reference effects and the explanation (formal definition in Def 3, Appendix). Following Abraham et al. (2022), we consider three distance metrics: cosine distance, L2 distance, and norm difference. We use their mean as the final reported error ($\overline{\text{ED}}$).

In addition, we propose a new measure, *Order-Faithfulness*, which builds on the necessary condition for faithful explanations introduced by Gat et al. (2023), which states that an explanation must rank one concept as more important than another if and only if its true causal effect is larger. While ED measures estimation accuracy, Order-Faithfulness assesses whether explanations preserve the relative ordering of concept importance, a property that is often more robust, interpretable, and directly relevant to how explanations are used in practice. To formalize this idea, consider two concept changes \vec{c}_1 and \vec{c}_2 . We first compute the difference between their reference effect vectors, and then the difference between their explanation vectors. We compare the signs of each entry in the difference vector with the corresponding entry in the explanation difference vector. Agreement of signs indicates that the explanation preserves the correct ordering of

³No training is required if f is a zero/few-shot LLM.

Box 3.1: Definitions: Evaluation Measures

Definition 3 (ICaCE Error Distance (ED)).

$$\begin{aligned} \text{ED}(f, M_f, x_\varepsilon, \vec{c}) &= \text{dist}\left(\widehat{\text{ICaCE}}_f(x_\varepsilon, \vec{c}); M_f(x_\varepsilon, \vec{c})\right) \\ \overline{\text{ED}}(f, M_f) &= \frac{1}{|\mathcal{C}|} \sum_{\vec{c} \in \mathcal{C}} \frac{1}{|D_{\vec{c}}|} \sum_{x_\varepsilon \in D_{\vec{c}}} \text{ED}(f, M_f, x_\varepsilon, \vec{c}) \end{aligned}$$

Definition 4 (ICaCE Order-Faithfulness (OF)).

$$\begin{aligned} \text{OF}(f, M_f, x_\varepsilon, \vec{c}_1, \vec{c}_2) &= \text{sign}\left(\widehat{\text{ICaCE}}_f(x_\varepsilon, \vec{c}_1) - \widehat{\text{ICaCE}}_f(x_\varepsilon, \vec{c}_2); M_f(x_\varepsilon, \vec{c}_1) - M_f(x_\varepsilon, \vec{c}_2)\right) \\ \overline{\text{OF}}(f, M_f) &= \frac{1}{|\mathcal{C}|(|\mathcal{C}| - 1)} \sum_{\substack{\vec{c}_1, \vec{c}_2 \in \mathcal{C} \\ \vec{c}_1 \neq \vec{c}_2}} \frac{1}{|D_{\vec{c}_1} \cap D_{\vec{c}_2}|} \sum_{x_\varepsilon \in D_{\vec{c}_1} \cap D_{\vec{c}_2}} \text{OF}(f, M_f, x_\varepsilon, \vec{c}_1, \vec{c}_2) \end{aligned}$$

Where $\text{dist}(\cdot; \cdot)$ is a distance metric and $\text{sign}(\cdot; \cdot)$ is the proportion of vector entries that agree in sign.

the changes, and is therefore order-faithful.

The formal definition is provided in Def 4 (Appendix). To summarize, we report the average error distance $\overline{\text{ED}}$ (lower is better) and the average order-faithfulness $\overline{\text{OF}}$ (higher is better).

4 Interventional Data Generation

We next describe the process for generating an interventional benchmark using LIBERTy (LLM-based Interventional Benchmark for Explainability with Reference Targets). The framework relies on explicitly defined DGPs comprising three components: SCMs over concepts, exogenous grounding texts, and an LLM (see Figure 1). These DGPs allow us to generate silver counterfactuals.

LIBERTy SCMs For each dataset, we first define the causal graph that specifies the concepts and their directional relationships (which concept influences which). Based on this graph, we specify the structural equations: each concept is linked to a function that determines its value based on its parent concepts and an exogenous noise term. The noise is drawn from a Gaussian distribution, with a concept-specific mean and variance. The structural equation generating the text takes all concepts as inputs and uses two exogenous grounding texts, a persona and a template, instead of Gaussian noise.

In Figure 2, we illustrate the three causal graphs of the three LIBERTy datasets. While their SCMs are not intended to mirror the true causal structure of the world (see the discussion in Appendix A.1), they are grounded in plausible assumptions: one causal graph (workplace violence prediction) is adapted from prior literature (Gerberich et al., 2004), and the other two (disease detection and CV screening) are informed by statistical patterns in real-world data (Monto et al., 2000; Cady and Schreiber, 2002; Dastin, 2018). Finally, we note

that our three causal graphs are much more complex and richer than the (four-concepts) causal graph of CEBaB. Each graph includes at least eight concepts, exhibits confounding and mediation structures (allowing estimation of direct and indirect effects), contains long paths (up to four edges between a concept and the text), and supports both anticausal ($Y \rightarrow \text{Text}$) and confounded ($Y \leftarrow C \rightarrow \text{Text}$) learning problems.

Exogenous Grounding Texts To ensure the validity of our structural counterfactuals, deterministic decoding is required. With stochastic decoding, generation noise cannot be tracked or held fixed across factual and counterfactual texts, causing them to differ in unobserved exogenous factors rather than only in the intervened concepts, and thus violating the definition of a structural counterfactual (see Appendix A.2). However, this requirement introduces its own limitations. First, for a given combination of concept values, deterministic decoding produces a single, fixed text. Second, this decoding yields highly generic, templated, and repetitive texts, regardless of concept values (always the same narrative, albeit with minor variations). Third, the generated examples do not seem like authentic human-written text. To address these limitations, we propose a simple yet elegant solution in the spirit of the SCM framework: we introduce two additional exogenous variables, an author persona and a template, both of which serve as a grounding context for the LLM.

The *Persona variable* $\varepsilon_{\text{persona}}$ represents a set of contextual attributes, including profession, hobbies, and personal motivations. In contrast, the *Template variable* $\varepsilon_{\text{template}}$ captures a particular discourse structure, derived from real-world corpora (e.g., personal statements, Reddit posts). Templates and personas support three key goals: (1) making the

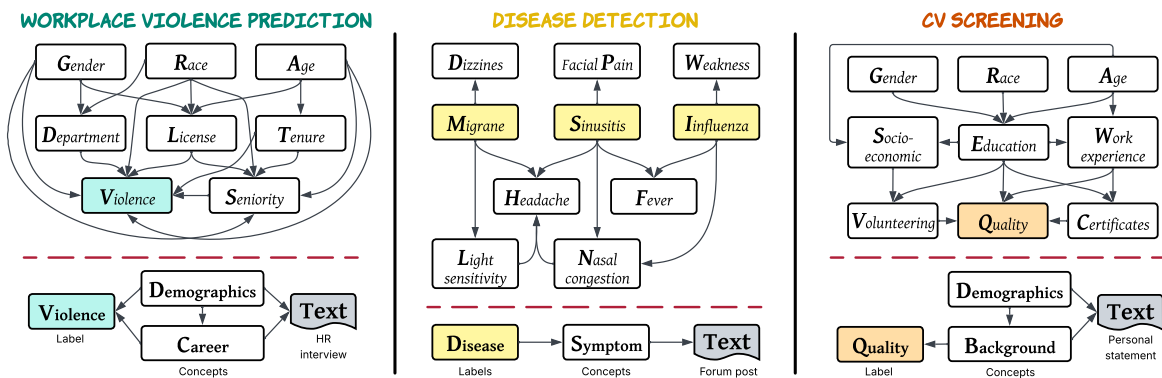


Figure 2: **LIBERTy Causal Graphs:** We show only the concepts (endogenous variables) and the relationships between them. Colored concepts indicate the variables that the explained model is trained to predict (the Y). At the bottom, a simplified version is provided. The graphs are grounded in prior literature and studies.

generated texts resemble authentic text; (2) promoting diversity: for each set of concept values, there are $|\mathcal{E}_{persona}| \times |\mathcal{E}_{template}|$ possible instantiations; and (3) ensuring the original example and its counterfactual derive from the same narrative.

Text Generation We sample concept values in topological order from the SCM, using the equations and Gaussian noise. We then sample a persona and a template and record all variable values for later counterfactual generation. Textual realizations are generated via deterministic decoding (zero temperature) by conditioning GPT-4o on the full set of concept values, along with the persona and template. We use a dedicated prompt for each dataset. Notably, GPT-4o receives only the concept values and does not observe the causal graph itself.

Counterfactual Generation We follow Pearl’s three-step counterfactual procedure (Pearl, 2009). (1) Abduction: fix the exogenous variables used for the original example, (2) Action: intervene on a target concept, and (3) Prediction: propagate the intervention through the SCM and compute updated concept values. We then regenerate the text using the same persona, template, and deterministic decoding. The red arrows in Figure 1 illustrate this. For each test example, we randomly select three concept changes and generate counterfactuals.

5 Datasets

LIBERTy comprises three datasets that cover high-stakes, socially impactful NLP tasks. Each dataset is generated from an explicit SCM and split into four subsets: train and test sets for the explained model, a set for training explanation methods, and a counterfactual test set containing factual-

counterfactual pairs. Dataset statistics appear in Table 6. The LLM integrated within the SCMs (for generating texts) is GPT-4o, while Gemini-1.5-Pro is used to create templates and personas. Due to space limitations, we only briefly describe each dataset here; full details of their construction, including motivations, complete SCMs, prompts, and generated examples, are provided in Appendix E. In Appendix B, we conduct a human evaluation of the quality of our generated data.

Workplace Violence Prediction: This dataset simulates HR–nurse interviews, in which the (explained) model predicts the likelihood that a nurse will experience workplace violence. The causal graph is adapted from the Minnesota Nurses’ Study (Gerberich et al., 2004) and captures demographic and professional risk factors. The template follows a structured interview, in which we sample one question per concept from a question bank and randomly order the questions. The persona consists of three informal “fun facts” about the nurse.

Disease Detection: This dataset simulates clinical self-report posts describing symptoms on online medical forums, in which the model predicts a disease. The SCM is anti-causal, with the disease as the root cause determining symptoms based on established symptom–disease relations (Monto et al., 2000; Cady and Schreiber, 2002). The template is a narrative abstracted by Gemini from 1,310 posts on Reddit’s DiagnoseMe forum. The persona (a total of 1200) consists of three informal facts about occupation, hobbies, and family or friends.

CV Screening: This dataset simulates automated resume screening, where the model predicts applicant quality from a CV-style personal statement.

The SCM encodes dependencies between demographic and professional attributes, inspired by statistical patterns reported by the U.S. Bureau of Labor Statistics. 1,235 templates were abstracted from 342 scraped personal statement examples,⁴. The persona (a total of 990) comprises role-specific personal and professional background details, including motivations and skills.

6 Experimental Setup

Using LIBERTy, we conduct experiments on five explained models and benchmark eight explanation methods from four families of approaches. The goals of our experiments are: (1) Benchmarking local and global explanation methods; (2) Analyzing the sensitivity of models to concept changes and evaluating which model captures better the causal structure of the data. The evaluation pipeline is described in Section 3.2. When reporting scores, we typically average them over all concept changes.

Explained Models We evaluate five models. Three are fine-tuned to predict Y from text: (1) DeBERTa-v3 (base, He et al. (2020)), an encoder-only model; (2) T5 (base, Raffel et al. (2020)), an encoder-decoder model; and (3) Qwen-2.5 (1.5B-instruct, Team (2023)), a decoder-only LLM. The other two are zero-shot LLMs: (4) Llama-3.1 (8B-instruct, Dubey et al. (2024)) and (5) GPT-4o (OpenAI, 2024). See Appendix F.2 for more details, hyperparameters, performance, and prompts.

Explainability Methods We briefly mention the explainability methods we benchmark, but Appendix D thoroughly describes and discusses them. The rationale for selecting methods was to focus on top-performing approaches previously applied to CEBaB with user-friendly code. We examine eight methods covering four families:

(1) *Counterfactual Generation*: LLMs generate counterfactuals by editing texts to reflect a target concept change (Gat et al., 2023). We examine in Appendix D.1 four prompting techniques, each injecting different causal assumptions. We mainly focus on the *Mediators and Confounders* technique, which fixes confounders while allowing mediators to vary, and achieves the best performance.

(2) *Matching* (see Appendix D.2): matching methods search for the most similar candidate from a predefined set of examples with the target concept change. The difference between the methods lies

in how similarity is defined, and we examine five methods: (2a) *ST Match*: cosine similarity over SentenceTransformer embeddings (Reimers and Gurevych, 2019); (2b) *PT Match*: cosine similarity over a pre-trained encoder-only model (DeBERTa); (2c) *FT Match*: cosine similarity over an encoder fine-tuned to predict Y ; (2d) *Approx*: first predicts concept values using fine-tuned models and then search for exact concept-based match; and (2e) *ConVecs*: cosine similarity over concatenated softmax prediction vectors of all concepts. Notably, the first three are semantic-based methods, while the latter two are concept-based ones.

(3) *Concept Erasure* (see Appendix D.3): removes linearly encoded information about a target concept from hidden representations using LEACE (Belrose et al., 2023).⁵ (4) *Concept Attributions* (see Appendix D.4): estimates concept importance via *ConceptShap* (Yeh et al., 2020) combined with *TCAV* (Kim et al., 2018), which construct concept vectors and assign Shapley-based scores.⁶

7 Results

7.1 Local Explanations

We begin by comparing the local explainability methods using LIBERTy, reporting ICaCE \overline{ED} and \overline{OF} . Table 1 presents these measures at the dataset level (averaged across all five models) and at the model level (averaged across all three datasets). Complete results are provided in Table 15 (Appendix G). Overall, the matching approach performs best. Within this category, *FT Match*, which fine-tunes an encoder-only model to predict the label Y and then uses its embeddings for similarity, achieves the lowest estimation error and emerges as the most faithful method. Its advantage likely stems from the model learning task-specific representations that produce more meaningful neighborhoods for matching. Other strong performers are the concept-based matching methods, *ConVecs* (proposed in this work) and *Approx*. These findings align with those of Gat et al. (2023), who compared different matching methods on CEBaB and reported similar trends.

An interesting difference between our findings and those of Gat et al. (2023) is that, while LLM-generated counterfactuals outperform matching-

⁵We employ LEACE only for open-source models and only on the Disease Detection dataset, where erasing a concept is well defined as its absence (e.g., symptom not present).

⁶We benchmark *ConceptShap* only as a global explanation for open-source models.

⁴<https://universitycompare.com>

↓ Method	Average		Dataset						Explained Model					
			Violence		Disease		CV		DeBERTa-v3		Qwen-2.5		GPT-4o	
	ED	OF	ED	OF	ED	OF	ED	OF	ED	OF	ED	OF	ED	OF
<i>CF Gen</i>	0.55	0.49	0.47	0.58	0.67	0.36	0.52	0.52	0.50	0.59	0.62	0.53	0.58	0.49
<i>Approx</i>	0.45	0.69	0.41	0.71	0.48	0.69	0.46	0.66	0.38	0.76	0.50	0.70	0.53	0.67
<i>ConVecs</i>	0.44	0.69	0.40	0.73	0.44	0.70	0.47	0.66	0.34	0.78	0.47	0.71	0.52	0.68
<i>ST Match</i>	0.49	0.65	0.51	0.63	0.46	0.69	0.50	0.62	0.49	0.69	0.55	0.66	0.53	0.67
<i>PT Match</i>	0.51	0.64	0.51	0.64	0.52	0.65	0.50	0.63	0.52	0.68	0.56	0.65	0.59	0.64
<i>FT Match</i>	0.34	0.74	0.32	0.76	0.36	0.75	0.35	0.72	0.16	0.88	0.39	0.75	0.48	0.70
<i>LEACE</i>	0.65	0.46	—	—	0.65	0.46	—	—	0.62	0.42	0.87	0.41	—	—

Table 1: **Local Explainability Results:** We report the Average ICaCE Error-Distance ($\overline{ED} \geq 0$, \downarrow is better) and Average ICaCE Order-Faithfulness ($\overline{OF} \leq 1$, \uparrow is better). The **Average** column reports the mean across five explained models and three datasets. The detailed results appear in Appendix Table 15 and exhibit a similar pattern, with fine-tuned matching outperforming other approaches. Horizontal lines separate method families.

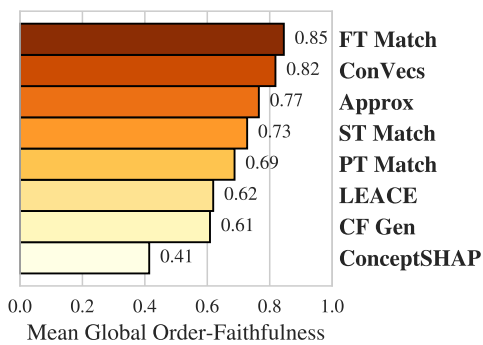


Figure 3: **Global Explainability Results:** We report the mean Order-Faithfulness score for global explanations. See Table 16 in the Appendix for full results.

based methods on CEBaB, the opposite holds on LIBERTY. A potential explanation is that humans write CEBaB’s counterfactuals: annotators edit an existing text to reflect a change in concept. LLMs can closely mimic this editing process, especially for short, simple texts, which makes their generated counterfactuals appear effective. In LIBERTY, producing an explanation that resembles a human edit does not guarantee faithfulness; instead, the explanation should reflect the actual DGP. This also explains why matching methods perform more consistently: their retrieved candidates are sampled from distributions aligned with the underlying DGP rather than produced through human-aligned textual edits. We refer the reader to an extended discussion of these aspects in Appendix A.3.

Finally, the \overline{ED} and \overline{OF} scores reveal substantial room for improvement. In LIBERTY, even the best methods achieve only around 0.3 on ED (where 0 is perfect) and 0.7 on OF (where 1 is perfect). We hope that LIBERTY will encourage further progress on developing more faithful explanation methods.

Dataset	Violence	Disease	CV
Model	Qwen-2.5	DeBERTa-v3	GPT-4o
Gold	Gender	Light Sens	Work Exp
	Department	Facial Pain	Education
	Age	Dizziness	Race
FT Match	Gender	Light Sens	Education
	Seniority	Dizziness	Work Exp
	Age	Facial Pain	Age
CF Gen	Gender	Weakness	Education
	Age	Dizziness	Work Exp
	Race	Light Sens	Socioeco
LEACE		Dizziness	
		Light Sens	
		Headache	
ConceptShap	Gender	Dizziness	
	Race	Nasal Cong	
	Seniority	Weakness	

Table 2: **Global Explanations Analysis:** We present the top-3 most important concepts of explanations for selected datasets, models, and methods. A colored concept indicates it is among the top three gold concepts.

7.2 Global Explanations

Many global explanations produce a ranked list of concepts by their overall importance (not specific to a single example), reflecting their influence on the model’s predictions (a.k.a. feature importance). We therefore evaluate their global order-faithfulness: are the concepts ranked in the same order as their causal effects? To obtain the ground-truth ranking, we compute a single gold importance score for each concept using CaCE. Note that for each concept change, CaCE yields a vector of size $|Y|$, capturing the causal effect of that change on each output class. To obtain a single gold importance score for each concept, we first sum the absolute CaCE values across all output classes, reflecting the total magnitude of the effect for that change. We then average this quantity over all changes. This produces a single gold importance

Examined Model	Workplace Violence			Disease Detection		CV Screening		
	Race	Gender	Age	Headache	General Weakness	Race	Gender	Age
DeBERTa-v3	0.350	1.192	0.758	0.398	0.415	0.715	0.432	0.613
T5	0.421	0.743	0.512	0.530	0.376	0.742	0.398	0.513
Qwen-2.5	0.691	1.314	1.045	0.426	0.512	0.522	0.361	0.503
Llama-3.1	0.224	0.227	0.226	0.364	0.332	0.374	0.283	0.397
GPT-4o	0.724	0.594	0.300	0.369	0.215	0.417	0.208	0.355
True Effect	0.484	1.271	1.154	–	–	0.636	0.369	0.913

Table 3: **Concept Sensitivity Analysis:** In the Disease Detection dataset, Y is the parent of the concepts, so interventions do not affect Y , and its ground-truth sensitivity cannot be computed. See Table 17 for full results.

score per concept. Global \overline{OF} is then computed based on these scores: it quantifies how faithfully each explanation method’s ranking of concept importance matches the gold ranking.⁷

Figure 3 compares the methods using the average global \overline{OF} across the three datasets and five models. The complete (non-averaged) results are reported in Table 16 in the Appendix. As shown, global trends mirror the local ones, with the matching approach outperforming the others. Table 2 further reports the top-3 most important concepts identified by each method and compares them with the top-3 gold concepts (according to their gold importance score). Every method misses at least one gold concept, highlighting the need for further research on global explainability.

7.3 Sensitivity Analysis

Up to this point, we have used LIBERTy to evaluate explanation methods. More broadly, the framework supports two complementary analyses. First, it can be used to analyze a model’s sensitivity to concept changes by measuring the magnitude of prediction changes induced by structural counterfactuals. Second, LIBERTy can be used to assess how well different learning methods, such as CE fine-tuning, align model behavior with the causal structure encoded in the DGP. This second analysis necessarily focuses on models trained on the generated data, since only then can their behavior be expected to reflect the underlying causal relationships. Under successful learning, a model’s sensitivity to concept changes should closely match the true causal effects on the outcome variable Y , which we estimate via Monte Carlo simulation from the SCM.

For a given example and concept change, we

⁷Global \overline{OF} and ICaCE \overline{OF} differ both in the order of computation and in what is being ranked. ICaCE \overline{OF} evaluates order-faithfulness over individual concept changes on a per-example basis before averaging, whereas Global \overline{OF} evaluates order-faithfulness over concepts, using global importance scores derived from CaCE.

compute a sensitivity score that quantifies the extent to which the model’s prediction is affected. This score is obtained by summing the absolute ICaCE values, which quantify the magnitude of the change. Larger values indicate stronger shifts in the prediction (i.e., more sensitive). Table 3 reports sensitivity scores for the five evaluated models on selected concepts (an average over all their changes), alongside the gold sensitivity effect. Complete results are provided in Table 17 in the Appendix.

When examining sensitivity scores (without comparing them to the gold effects) we observe that zero-shot LLMs (Llama-3.1-8B and GPT-4o) exhibit lower sensitivity to concept changes, particularly for demographic concepts such as Race, Gender, and Age (Table 3). We believe the reduced sensitivity reflects intentional design choices made during post-training alignment. In addition, among the fine-tuned models, we find that Qwen2.5-1.5B most accurately reflects the causal structure of the data. Still, the gap with the gold effects highlights that fine-tuning is insufficient and that there remains a need for causal learning techniques.

8 Conclusions

A central challenge in explainability is the lack of reliable evaluation protocols, particularly given the absence of “gold explanations”. Our work takes a significant step toward closing this gap. We introduced LIBERTy, a framework for generating interventional datasets to benchmark concept-based explanations against “silver” references estimated using structural counterfactuals. In Section A.4 in the Appendix, we outline future research opportunities motivated by our key findings. To summarize, there is great promise in developing smaller, theory-grounded, causal-inspired explainability and learning approaches. We hope our work will serve as a foundation for such future research.

9 Limitations

Synthetic Text Generation LIBERTy relies on LLMs to instantiate structural counterfactuals. However, it also means that the texts are synthetic rather than human-written. This may introduce mismatches between how the LLM instantiates concepts and how humans would naturally express them. To assess data quality, we conducted a human evaluation (Appendix B). Annotators confirmed that the generated texts are coherent, relevant, and fluent; that the LLM correctly incorporates concept values; and that counterfactuals are perceived as realistic variants differing in only one concept. Finally, although LIBERTy uses synthetic text, this limitation is increasingly less restrictive: a growing share of real-world data is generated by LLMs, making synthetic settings both common and practically meaningful. It is therefore reasonable to assume that model inputs in many future applications will themselves be LLM-generated.

Focusing on Concept-based Explanations Our work focuses exclusively on concept-based explanations and their causal evaluation. This scope covers only a subset of existing explainability methods, and most prior work centers on token-level or free-text explanations (see the analysis of Calderon and Reichart (2024)). Nevertheless, there are strong reasons to focus on concept-based methods. These methods quantify how high-level, human-interpretable concepts (a.k.a. attributes, features, variables, or rubrics) that are implicitly or explicitly expressed in text influence the model. Because such high-level concepts align with human cognitive processes (Alqaraawi et al., 2020; Kim et al., 2022; Poeta et al., 2023), reduce the complexity of long inputs, and communicate model behavior in intuitive terms (Calderon and Reichart, 2024), concept-based explanations are particularly suitable for high-stakes settings where end users and decision makers must understand and trust model reasoning. We believe that the relatively limited attention to concept-based explainability stems partly from the lack of appropriate benchmarks for developing and evaluating such methods. By providing an interventional benchmark with structural causal effects, LIBERTy aims to address this gap and facilitate broader research and adoption of concept-based explanations.

DGPs as Approximations of Reality LIBERTy provides structural counterfactuals in the strict

sense, as they are generated from a fully specified data-generating process (DGP). While the DGP and causal graphs only simplify real-world mechanisms, they are not arbitrary and are grounded in domain knowledge and the literature. Still, we acknowledge that they do not perfectly mirror real-world causal structures. Crucially, this limitation does not compromise the reliability of our evaluation protocol, because our goal is not to recover real-world mechanisms or estimate real-world causal effects. Instead, our objective is to measure the causal effects *within the explained model* and benchmark explanation methods against those effects. For this purpose, what matters is that the DGP supports precise interventions and produces structural counterfactuals that faithfully reflect them. Explanation faithfulness is always defined relative to the explained model, whether its behavior arises from true causal relationships, simplified abstractions, or even spurious correlations. Thus, a synthetic DGP is sufficient and, in practice, often required for the controlled and rigorous evaluation of explanation methods. Such methods can be trained on or applied to data generated by the DGP and evaluated against the explained model’s predictions, whether or not the model itself was trained on that data. We do not claim that our benchmark explains real-world phenomena or reveals how LLMs internally represent them. Rather, our goal is to provide a principled benchmark for comparing explanation methods, analyzing their limitations, and identifying those that most faithfully capture model behavior, thereby enabling their application in real-world settings. Please also see our discussion in Appendix A.1.

Acknowledgments

This work has been partially funded by the Israel MOST grant on Improving human-context reasoning in vision-language models.

References

- Eldar David Abraham, Karel D’Oosterlinck, Amir Feder, Yair Ori Gat, Atticus Geiger, Christopher Potts, Roi Reichart, and Zhengxuan Wu. 2022. [Cebab: Estimating the causal effects of real-world concepts on NLP model behavior](#). *CoRR*, abs/2205.14140.
- Ahmed Alqaraawi, Martin Schuessler, Philipp Weiß, Enrico Costanza, and Nadia Berthouze. 2020. [Evaluating saliency map explanations for convolutional neural networks: A user study](#). *CoRR*, abs/2002.00772.

- Esma Balkir, Svetlana Kiritchenko, Isar Nejadgholi, and Kathleen C. Fraser. 2022. [Challenges in applying explainability methods to improve the fairness of NLP models](#). *CoRR*, abs/2206.03945.
- Nora Belrose, David Schneider-Joseph, Shauli Ravfogel, Ryan Cotterell, Edward Raff, and Stella Biderman. 2023. [LEACE: perfect linear concept erasure in closed form](#). *CoRR*, abs/2306.03819.
- Kenza Benkirane, Jackie Kay, and María Pérez-Ortiz. 2024. [How can we diagnose and treat bias in large language models for clinical decision-making?](#) *CoRR*, abs/2410.16574.
- Diane Bouchacourt and Ludovic Denoyer. 2019. [EDUCE: explaining model decisions through unsupervised concepts extraction](#). *CoRR*, abs/1905.11852.
- Roger K Cady and Curtis P Schreiber. 2002. [Sinus headache or migraine? considerations in making a differential diagnosis](#). *Neurology*, 58(9_suppl_6):S10–S14.
- Nitay Calderon, Liat Ein-Dor, and Roi Reichart. 2025. [Multi-domain explainability of preferences](#). *CoRR*, abs/2505.20088.
- Nitay Calderon and Roi Reichart. 2024. [On behalf of the stakeholders: Trends in NLP model interpretability in the era of llms](#). *CoRR*, abs/2407.19200.
- Fateme Hashemi Chaleshtori, Atreya Ghosal, Alexander Gill, Purbid Bamroo, and Ana Marasovic. 2024. [On evaluating explanation utility for human-ai decision making in NLP](#). *CoRR*, abs/2407.03545.
- Ivi Chatzi, Nina L. Corvelo Benz, Eleni Straitouri, Stratis Tsirtsis, and Manuel Gomez-Rodriguez. 2025. [Counterfactual token generation in large language models](#). In *Causal Learning and Reasoning, Lausanne, Switzerland, 7-9 May 2025*, volume 275 of *Proceedings of Machine Learning Research*, pages 1291–1315. PMLR.
- Giandomenico Cornacchia, Vito Walter Anelli, Fedelucio Narducci, Azzurra Ragone, and Eugenio Di Sciascio. 2023. [Counterfactual reasoning for bias evaluation and detection in a fairness under unawareness setting](#). *CoRR*, abs/2302.08204.
- Bo Cowgill, Fabrizio Dell’Acqua, Samuel Deng, Daniel Hsu, Nakul Verma, and Augustin Chaintreau. 2020. [Biased programmers? or biased data? A field experiment in operationalizing AI ethics](#). *CoRR*, abs/2012.02394.
- Fahim Dalvi, Abdul Rafae Khan, Firoj Alam, Nadir Durrani, Jia Xu, and Hassan Sajjad. 2022. [Discovering latent concepts learned in BERT](#). *CoRR*, abs/2205.07237.
- Jeffrey Dastin. 2018. [Amazon scrapped ‘ai’ recruiting tool that showed bias against women](#).
- Jin Du, Li Chen, Xun Xian, An Luo, Fangqiao Tian, Ganghua Wang, Charles Doss, Xiaotong Shen, and Jie Ding. 2025. [Ice cream doesn’t cause drowning: Benchmarking llms against statistical pitfalls in causal inference](#). *CoRR*, abs/2505.13770.
- Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, Anirudh Goyal, Anthony Hartshorn, Aobo Yang, Archi Mitra, Archie Sravankumar, Artem Korenev, Arthur Hinsvark, Arun Rao, Aston Zhang, and 82 others. 2024. [The llama 3 herd of models](#). *CoRR*, abs/2407.21783.
- Amir Feder, Nadav Oved, Uri Shalit, and Roi Reichart. 2021. [Causalm: Causal model explanation through counterfactual language models](#). *Comput. Linguistics*, 47(2):333–386.
- Yair Ori Gat, Nitay Calderon, Amir Feder, Alexander Chapanin, Amit Sharma, and Roi Reichart. 2023. [Faithful explanations of black-box NLP models using llm-generated counterfactuals](#). *CoRR*, abs/2310.00603.
- S Gerberich, T Church, P McGovern, and et al. 2004. [An epidemiological study of the magnitude and consequences of work related violence: the minnesota nurses’ study](#). *Occupational and Environmental Medicine*, 61(6):495–503.
- Amirata Ghorbani, James Wexler, James Y. Zou, and Been Kim. 2019. [Towards automatic concept-based explanations](#). In *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pages 9273–9282.
- Yash Goyal, Uri Shalit, and Been Kim. 2019. [Explaining classifiers with causal concept effect \(cace\)](#). *CoRR*, abs/1907.07165.
- Riccardo Guidotti, Anna Monreale, Franco Turini, Dino Pedreschi, and Fosca Giannotti. 2018. [A survey of methods for explaining black box models](#). *CoRR*, abs/1802.01933.
- Sai Gurrupu, Ajay Kulkarni, Lifu Huang, Ismini Lourentzou, and Feras A. Batarseh. 2023. [Rationalization for explainable NLP: a survey](#). *Frontiers Artif. Intell.*, 6.
- Peter Hase and Mohit Bansal. 2020. [Evaluating explainable AI: which algorithmic explanations help users predict model behavior?](#) *CoRR*, abs/2005.01831.
- Pengcheng He, Xiaodong Liu, Jianfeng Gao, and Weizhu Chen. 2020. [Deberta: Decoding-enhanced BERT with disentangled attention](#). *CoRR*, abs/2006.03654.
- Anna Hedström, Philine Lou Bommer, Kristoffer Knutson Wickström, Wojciech Samek, Sebastian Lapuschkin, and Marina M.-C. Höhne. 2023. [The meta-evaluation problem in explainable AI: identifying](#)

- reliable estimators with metaquantus. *Trans. Mach. Learn. Res.*, 2023.
- XinYue Jiang, Jingsong He, and Li Gu. 2025. [MTCR: method for matching texts against causal relationship](#). *Neural Process. Lett.*, 57(3):58.
- Enkelejda Kasneci, Kathrin Sessler, Stefan Küchermann, Maria Bannert, Daryna Dementieva, Frank Fischer, Urs Gasser, Georg Groh, Stephan Günemann, Eyke Hüllermeier, Stephan Krusche, Gitta Kutyniok, Tilman Michaeli, Claudia Nerdel, Jürgen Pfeffer, Oleksandra Poquet, Michael Sailer, Albrecht Schmidt, Tina Seidel, and 2 others. 2023. [Chatgpt for good? on opportunities and challenges of large language models for education](#). *ScienceDirect*.
- Been Kim, Martin Wattenberg, Justin Gilmer, Carrie J. Cai, James Wexler, Fernanda B. Viégas, and Rory Sayres. 2018. [Interpretability beyond feature attribution: Quantitative testing with concept activation vectors \(TCAV\)](#). In *Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholmsmässan, Stockholm, Sweden, July 10-15, 2018*, volume 80 of *Proceedings of Machine Learning Research*, pages 2673–2682. PMLR.
- Sunnie S. Y. Kim, Elizabeth Anne Watkins, Olga Russakovsky, Ruth Fong, and Andrés Monroy-Hernández. 2022. ["help me help the ai": Understanding how explainability can support human-ai interaction](#). *CoRR*, abs/2210.03735.
- Pang Wei Koh, Thao Nguyen, Yew Siang Tang, Stephen Mussmann, Emma Pierson, Been Kim, and Percy Liang. 2020. [Concept bottleneck models](#). In *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event*, volume 119 of *Proceedings of Machine Learning Research*, pages 5338–5348. PMLR.
- Sören R. Künnel, Jasjeet S. Sekhon, Peter J. Bickel, and Bin Yu. 2019. [Meta-learners for estimating heterogeneous treatment effects using machine learning](#). *arXiv*.
- Jun Rui Lee, Sadegh Emami, Michael David Hollins, Timothy C. H. Wong, Carlos Ignacio Villalobos Sánchez, Francesca Toni, Dekai Zhang, and Adam Dejl. 2025. [Xai-units: Benchmarking explainability methods with unit tests](#). *CoRR*, abs/2506.01059.
- Yongqi Li, Mayi Xu, Xin Miao, Shen Zhou, and Tiejun Qian. 2024. [Prompting large language models for counterfactual generation: An empirical study](#). In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation, LREC/COLING 2024, 20-25 May, 2024, Torino, Italy*, pages 13201–13221. ELRA and ICCL.
- Scott M. Lundberg and Su-In Lee. 2017. [A unified approach to interpreting model predictions](#). *CoRR*, abs/1705.07874.
- Siwen Luo, Hamish Ivison, Soyeon Caren Han, and Josiah Poon. 2024. [Local interpretations for explainable natural language processing: A survey](#). *ACM Comput. Surv.*, 56(9):232:1–232:36.
- Qing Lyu, Marianna Apidianaki, and Chris Callison-Burch. 2022. [Towards faithful model explanation in NLP: A survey](#). *CoRR*, abs/2209.11326.
- Arnold S Monto, Stefan Gravenstein, Michael Elliott, Michael Colopy, and Jo Schweinle. 2000. [Clinical signs and symptoms predicting influenza infection](#). *Archives of internal medicine*, 160(21):3243–3247.
- Yuqi Nie, Yaxuan Kong, Xiaowen Dong, John M. Mulvey, H. Vincent Poor, Qingsong Wen, and Stefan Zohren. 2024. [A survey of large language models for financial applications: Progress, prospects and challenges](#). *CoRR*, abs/2406.11903.
- OpenAI. 2024. [Gpt-4o technical report](#). *arXiv preprint*.
- Judea Pearl. 2009. *Causality: Models, Reasoning, and Inference*, 2 edition. Cambridge University Press.
- Judea Pearl. 2013. [Structural counterfactuals: A brief introduction](#). *Cognitive science*, 37(6):977–985.
- Lotem Peled-Cohen, Maya Zadok, Nitay Calderon, Hila Gonen, and Roi Reichart. 2025. [Dementia through different eyes: Explainable modeling of human and LLM perceptions for early awareness](#). *CoRR*, abs/2505.13418.
- Eleonora Poeta, Gabriele Ciravegna, Eliana Pastor, Tania Cerquitelli, and Elena Baralis. 2023. [Concept-based explainable artificial intelligence: A survey](#). *CoRR*, abs/2312.12936.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2020. [Exploring the limits of transfer learning with a unified text-to-text transformer](#). *J. Mach. Learn. Res.*, 21:140:1–140:67.
- Manish Raghavan, Solon Barocas, Jon M. Kleinberg, and Karen Levy. 2020. [Mitigating bias in algorithmic hiring: evaluating claims and practices](#). In *FAT* '20: Conference on Fairness, Accountability, and Transparency, Barcelona, Spain, January 27-30, 2020*, pages 469–481. ACM.
- Shauli Ravfogel, Anej Svete, Vésteinn Snæbjarnarson, and Ryan Cotterell. 2025. [Gumbel counterfactual generation from language models](#). In *The Thirteenth International Conference on Learning Representations, ICLR 2025, Singapore, April 24-28, 2025*. OpenReview.net.
- Shauli Ravfogel, Michael Twiton, Yoav Goldberg, and Ryan Cotterell. 2022. [Linear adversarial concept erasure](#). *CoRR*, abs/2201.12091.
- Nils Reimers and Iryna Gurevych. 2019. [Sentence-bert: Sentence embeddings using siamese bert-networks](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and*

- the 9th International Joint Conference on Natural Language Processing, EMNLP-IJCNLP 2019, Hong Kong, China, November 3-7, 2019*, pages 3980–3990. Association for Computational Linguistics.
- Marco Túlio Ribeiro, Sameer Singh, and Carlos Guestrin. 2016. "why should I trust you?": Explaining the predictions of any classifier. *CoRR*, abs/1602.04938.
- Marcel Robeer, Floris Bex, and Ad Feelders. 2021. Generating realistic natural language counterfactuals. In *Findings of the Association for Computational Linguistics: EMNLP 2021, Virtual Event / Punta Cana, Dominican Republic, 16-20 November, 2021*, pages 3611–3625. Association for Computational Linguistics.
- Wojciech Samek, Grégoire Montavon, Sebastian Lapuschkin, Christopher J. Anders, and Klaus-Robert Müller. 2021. Explaining deep neural networks and beyond: A review of methods and applications. *Proc. IEEE*, 109(3):247–278.
- Pratinav Seth and Vinay Kumar Sankarapu. 2025. Bridging the gap in xai-why reliable metrics matter for explainability and compliance. *CoRR*, abs/2502.04695.
- Ruihao Shui, Yixin Cao, Xiang Wang, and Tat-Seng Chua. 2023. A comprehensive evaluation of large language models on legal judgment prediction. *CoRR*, abs/2310.11761.
- Qwen Team. 2023. Qwen: The official repo of qwen chat.
- James Thorne, Andreas Vlachos, Christos Christodoulopoulos, and Arpit Mittal. 2019. Generating token-level explanations for natural language inference. *CoRR*, abs/1904.10717.
- Victor Veitch, Dhanya Sridhar, and David M. Blei. 2020. Adapting text embeddings for causal inference. In *Proceedings of the Thirty-Sixth Conference on Uncertainty in Artificial Intelligence, UAI 2020, virtual online, August 3-6, 2020*, volume 124 of *Proceedings of Machine Learning Research*, pages 919–928. AUAI Press.
- Lijie Wang, Yaozong Shen, Shuyuan Peng, Shuai Zhang, Xinyan Xiao, Hao Liu, Hongxuan Tang, Ying Chen, Hua Wu, and Haifeng Wang. 2022. A fine-grained interpretability evaluation benchmark for neural NLP. *CoRR*, abs/2205.11097.
- Yongjie Wang, Xiaoqi Qiu, Yu Yue, Xu Guo, Zhiwei Zeng, Yuhong Feng, and Zhiqi Shen. 2024. A survey on natural language counterfactual generation. In *Findings of the Association for Computational Linguistics: EMNLP 2024, Miami, Florida, USA, November 12-16, 2024*, pages 4798–4818. Association for Computational Linguistics.
- Zach Wood-Doughty, Ilya Shpitser, and Mark Dredze. 2018. Challenges of using text classifiers for causal inference. *CoRR*, abs/1810.00956.
- Tongshuang Wu, Marco Túlio Ribeiro, Jeffrey Heer, and Daniel S. Weld. 2021. Polyjuice: Generating counterfactuals for explaining, evaluating, and improving models. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing, ACL/IJCNLP 2021, (Volume 1: Long Papers), Virtual Event, August 1-6, 2021*, pages 6707–6723. Association for Computational Linguistics.
- Zhengxuan Wu, Karel D’Oosterlinck, Atticus Geiger, Amir Zur, and Christopher Potts. 2022. Causal proxy models for concept-based model explanations. *CoRR*, abs/2209.14279.
- Zhengxuan Wu, Atticus Geiger, Christopher Potts, and Noah D. Goodman. 2023. Interpretability at scale: Identifying causal mechanisms in alpaca. *CoRR*, abs/2305.08809.
- Fan Yang, Mengnan Du, and Xia Hu. 2019. Evaluating explanation without ground truth in interpretable machine learning. *CoRR*, abs/1907.06831.
- Chih-Kuan Yeh, Been Kim, Sercan Ömer Arik, Chun-Liang Li, Tomas Pfister, and Pradeep Ravikummar. 2020. On completeness-aware concept-based explanations in deep neural networks. In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*.
- Wei Jie Yeo, Ranjan Satapathy, and Erik Cambria. 2024. Towards faithful natural language explanations: A study using activation patching in large language models. *CoRR*, abs/2410.14155.
- Xuemin Yu, Fahim Dalvi, Nadir Durrani, and Hassan Sajjad. 2024. Latent concept-based explanation of NLP models. *CoRR*, abs/2404.12545.
- Raymond Zhang, Neha Nayak Kennard, Daniel Scott Smith, Daniel A. McFarland, Andrew McCallum, and Katherine Keith. 2023. Causal matching with text embeddings: A case study in estimating the causal effects of peer review policies. In *Findings of the Association for Computational Linguistics: ACL 2023, Toronto, Canada, July 9-14, 2023*, pages 1284–1297. Association for Computational Linguistics.
- Haiyan Zhao, Hanjie Chen, Fan Yang, Ninghao Liu, Huiqi Deng, Hengyi Cai, Shuaiqiang Wang, Dawei Yin, and Mengnan Du. 2024. Explainability for large language models: A survey. *ACM Trans. Intell. Syst. Technol.*, 15(2):20:1–20:38.

Appendix

A Discussion	14
A.1 Real-World Data	14
A.2 Deterministic Decoding	14
A.3 LLM-generated Counterfactuals	14
A.4 Opportunities	15
B Human Validation	15
C Extended Related Work	15
D Explainability Methods	16
D.1 Counterfactual Generation	17
D.2 Matching	17
D.3 Concept Erasure	18
D.4 Concept Attributions	18
E Dataset Details	19
E.1 Workplace Violence	19
E.2 Disease Detection	22
E.3 CV Screening	25
F Implementation Details	29
F.1 Explainability Methods	29
F.2 Explained Models	29
F.3 Prompts	29
G Additional Results	30

A Discussion

A.1 Real-World Data

Why is it acceptable that the LIBERTy SCM does not perfectly reflect the real world? We do not claim that our benchmark explains real-world phenomena or reveals how LLMs internally represent them. Rather, our goal is to provide a principled benchmark for comparing explanation methods, analyzing their limitations, and identifying those that most faithfully capture model behavior, thereby enabling their application in real-world settings. Therefore, it does not matter whether LIBERTy SCMs reflect real-world mechanisms (or are just inspired by them). Explainability faithfulness is defined with respect to the explained model, and an explanation method should account for the effects of concepts *as they are encoded by the model*, regardless of whether the model learns and represents real causal structures, synthetic structures, or spurious correlations.

A.2 Deterministic Decoding

Why is deterministic decoding necessary? While deterministic decoding has clear drawbacks, it is essential for LIBERTy. Counterfactual generation requires fixing the exogenous variables of the DGP. Yet, stochastic decoding introduces noise at the token-sampling level that lies outside the DGP and cannot be controlled or recorded. As a result, such counterfactuals cannot serve as ‘structural’ ones. Furthermore, they also fail by intuitive standards. Although the prompt for generating the original example and the counterfactual may differ only in one concept value, stochastic decoding often produces an entirely new narrative with little lexical overlap. While lexical overlap is not formally required, it remains a widely used proxy for counterfactual quality in NLP. Accordingly, many works generate counterfactuals by instructing LLMs to edit the original text minimally (Gat et al., 2023; Li et al., 2024; Wang et al., 2024). However, such examples are only approximations, since entirely different DGPs produce the original and counterfactual texts. Alternative solutions, beyond our approach of using exogenous grounding texts, include generating multiple counterfactuals and estimating ICaCE by averaging over them, or employing controlled decoding methods for counterfactual generation (Chatzi et al., 2025; Ravfogel et al., 2025).

A.3 LLM-generated Counterfactuals

Why do explanations based on LLM-generated counterfactuals fail? Explanations based on LLM-generated counterfactuals perform surprisingly well in benchmarks such as CEBaB (Gat et al., 2023), where human annotators provide the reference counterfactuals against which explanations are evaluated. However, this performance stems from the fact that both humans and LLMs approach the task similarly, by minimally editing the input text to reflect a change in concept. In such settings, LLMs can closely mimic the references, particularly when the texts are short and simple. Evaluation using human-written counterfactuals is therefore not an assessment of causal effects, but rather an evaluation of how well models mimic human editing. When evaluated under LIBERTy, however, the limitations of this approach become clear. Unlike human-written counterfactuals, LIBERTy provides structural counterfactuals derived from causal interventions in the DGP. LLM-generated counterfactuals fail in this setting

because their edits reflect heuristic assumptions, rather than the actual underlying mechanism.

A.4 Opportunities

What are the opportunities in the intersection between causality, explainability, and NLP/LLMs? Our findings reveal several exciting opportunities at the intersection of causality, explainability, and NLP. First, we observed a large room for improvement in both local and global explanations, offering clear targets for future work. There is clear potential for the development of causal-inspired explanation methods. Instead of relying on LLM-based explanations, which, despite encoding broad knowledge in their parameters, are not exposed to data from the target DGP and therefore fail to provide faithful explanations, small but principled techniques offer a more promising direction. These approaches can rely on causal structure rather than scale, making them especially well-suited for academic research. LIBERTy provides a rigorous evaluation ground for such methods and, we hope, will foster their further development.

Finally, our analysis revealed that vanilla fine-tuning may fail to capture the causal structure of the data, suggesting the need for unique learning methods. This opens an opportunity to harness LIBERTy as a testbed for developing and benchmarking new causal learning methods that go beyond fine-tuning, approaches that explicitly aim to align models with the underlying DGP. To summarize, there is great promise in developing smaller, theory-grounded, causal-inspired explainability and learning approaches. We hope our work will serve as a foundation for such future research.

B Human Validation

We conduct human validation of the generated examples to ensure: (1) they include all concept values; (2) they have high linguistic quality, by measuring coherence and fluency; (3) they are relevant to the task (e.g., look like a personal statement); (4) they are logically consistent with themselves and external facts; (5) the counterfactual feels like a genuine counterfactual (by measuring how likely the text was written by the same person in a parallel world where the concept value is different). Notably, human validation is not required to ensure that the LIBERTy evaluation pipeline is faithful; however, it helps demonstrate that the synthetically generated data is realistic and practical.

We recruited 13 annotators (all graduates with fluent English; 3 males, 10 females) who annotated a total of 349 single-text and 312 text-cf-pair evaluations. Each text was rated across six dimensions: five individual attributes assessing text-level quality and one comparative attribute assessing the quality of the counterfactual relative to its original. This resulted in a total of $349 \times 5 + 312 \times 1 = 2,057$ labels. The average inter-annotator agreement (IAA) across all dimensions is 0.91. The annotation guidelines can be viewed in Figures 4 and 5.

The results are presented in Table 4. As shown, the generated examples exhibit high linguistic quality, with average scores of 4.79 and 4.85 out of 5 for coherence and fluency, respectively. Their average scores for task relevance and logical consistency are 4.77 and 4.92. In addition, agreement with concept values is 94.2% on average, indicating that GPT-4o accurately instantiates the sampled values. The lowest scores appeared in the CV Screening dataset, probably because it involves socially sensitive concepts that are more heavily filtered during generation. Finally, annotators judged the counterfactuals to be genuine, with an average score of 4.44 out of 5, demonstrating that they were perceived as plausible even by the human eye.

C Extended Related Work

Concept-based Explainability Concept-based explainability encompasses methods that quantify the extent to which high-level, human-interpretable concepts (features, attributes, variables, rubrics) that can be explicitly or implicitly conveyed in the text influence model predictions. This is in contrast to token-level explanations, which emphasize tokens through techniques such as attribution or attention scores (Calderon and Reichart, 2024; Luo et al., 2024; Zhao et al., 2024). Concept-based explanations naturally align with human cognitive processes (Alqaraawi et al., 2020; Kim et al., 2022; Poeta et al., 2023) and simplify the complexity inherent in lengthy textual inputs, making explanations more intuitive and easier to communicate (Calderon and Reichart, 2024). Moreover, they naturally support both local and global explanations. Moreover, they naturally support two complementary explanation types: *local explanations*, which explain a single instance, and *global explanations*, which explain the entire data distribution (Gat et al., 2023). These advantages have driven their widespread use in applications such

	Workplace Violence	Disease Detection	CV Screening	Avg.
# Annotators	6	5	6	5.67
# Individual	76	170	103	116.33
# Pairs	101	105	106	104
# Labels	481	955	621	685.67
Avg. IAA	0.90	0.92	0.91	0.91
Avg. MAE	0.35	0.53	0.62	0.50
Concepts	97.9%	100%	84.7%	94.2%
Coherence	4.75	4.88	4.75	4.79
Fluency	4.72	4.90	4.92	4.85
Relevancy	4.80	4.68	4.83	4.77
Consistency	4.92	4.92	4.92	4.92
Plausibility	4.63	4.62	4.07	4.44

Table 4: **Results of Human Validation:** Average IAA and MAE are computed across annotator pairs: IAA for the binary concept identification task, and MAE for all other tasks using a 1–5 Likert scale. ‘Concepts’ reports the percentage of concept values that were marked as explicitly stated or logically inferred. ‘Plausibility’ reports the average score for a pair of texts being judged as an original and its counterfactual.

as bias detection (Cornacchia et al., 2023), providing clear and actionable explanations (Bouchacourt and Denoyer, 2019), discovering new hidden concepts (Ghorbani et al., 2019), explaining human preferences, reward models, and LLM-as-Judges (Calderon et al., 2025), and detecting dementia (Peled-Cohen et al., 2025).

Among the most prominent approaches of concept-based explainability, there are *Attribution methods* (Ribeiro et al., 2016; Lundberg and Lee, 2017; Kim et al., 2018; Yeh et al., 2020), *Concept Erasure methods* (Ravfogel et al., 2022; Belrose et al., 2023), *Counterfactual Generation methods*, (Feder et al., 2021; Robeer et al., 2021; Wu et al., 2021; Gat et al., 2023), and *Matching methods*, (Veitch et al., 2020; Zhang et al., 2023; Gat et al., 2023; Jiang et al., 2025), and *Concept Bottleneck models* (Koh et al., 2020; Dalvi et al., 2022; Yu et al., 2024). Using LIBERTy, we evaluate representative methods from those approaches. Nevertheless, despite the advantages of concept-based explainability, particularly for end-users and decision-makers, it remains underexplored relative to token-level approaches (Calderon and Reichart, 2024). A possible reason for this gap is the current lack of benchmarks that enable rigorous evaluation and systematic comparison.

Explainability Benchmarks Benchmarking explanations is a highly challenging task, primarily because ground-truth explanations are rarely available in real-world datasets (Yang et al., 2019; Hedström et al., 2023; Lee et al., 2025; Seth and

Sankarapu, 2025). Most prior evaluation methods relied on indirect proxies, such as checking whether different methods agree with one another or whether their outputs align with simple heuristics (Hase and Bansal, 2020; Samek et al., 2021). Furthermore, most explainability evaluations have focused on token-level explanations rather than reasoning over high-level semantic concepts (Thorne et al., 2019; Wang et al., 2022; Gurrapu et al., 2023). As mentioned earlier, *CEBaB* (Concept Effect Benchmark for NLP, Abraham et al. (2022)) was the first dataset to evaluate explainability methods under controlled interventions. CEBaB revealed that many popular methods fail to estimate causal effects accurately and often perform no better than a naive concept-based matching baseline.

Chaleshtori et al. (2024) recently noted an increasing need for richer benchmarks that capture the structural complexity of real-world data and enable the evaluation of both direct and indirect causal effects. Complementing this, Du et al. (2025) demonstrated that even state-of-the-art LLMs frequently fall prey to classical statistical fallacies, underscoring the limitations of existing evaluation methods in assessing true causal reasoning. We believe LIBERTy addresses these gaps by simulating realistic scenarios with diverse text types and rich causal graphs.

D Explainability Methods

In this section, we provide additional background on the explainability methods used in our study, as

well as further implementation details for each.

D.1 Counterfactual Generation

This approach uses an LLM (or a fine-tuned, pre-trained model when parallel training data are available) to generate approximations of counterfactuals. Typically, the LLM is instructed to modify the input text by replacing a specified concept with a target value. Gat et al. (2023) propose injecting causal assumptions into the prompt, in particular identifying confounder concepts from the causal graph and prompting the LLM to keep them fixed while changing the target concept. They found that LLM-generated counterfactuals yielded the best explanation method on CEBaB. In light of this, we extend their approach and compare different prompting strategies, each of which injects distinct causal assumptions into the prompt. In our causal graphs, relative to the target concept being modified, other concepts may play two key roles. The first are confounders, which act as root causes that influence both the target concept and the text, and therefore must remain fixed. The second are mediators, which are influenced by the target concept and, in turn, influence the text. They must be allowed to vary when measuring total causal effects.

The prompting techniques we evaluate are: (a) *Only Change*: specifies only the target concept change; (b) *Fix All*: specifies the change and instructs the LLM to fix the values of all other concepts; (c) *Fix Confounders*: specifies the change and the causal parents, explicitly forbidding their alteration.; (d) *Mediators and Confounders*: specifies all mediator concepts (without asking to fix their values) and the change, while instructing the LLM to fix the values of the confounding concepts.

To generate counterfactuals, we use Gemini-1.5-Pro, which differs from the LLM used to generate LIBERTy examples (GPT-4o). Importantly, although the prompts may mention the concepts and sometimes their roles (confounders or mediators), Gemini is expected to infer on its own how a change in the target concept affects other concepts (if they are mediators) and the resulting text. To compare different prompting techniques and manage computational costs, we restrict our experiments to the CV Screening dataset and three fine-tuned models: DeBERTa-base, T5-base, and Qwen2.5-1.5B. The results are reported in Table 5. As shown, the best-performing prompting technique is Mediators and Confounders, which is also the most causally informed. This technique explic-

itly incorporates both causal roles: it asks to hold the confounders fixed while allowing mediators to vary according to Gemini’s decision. Since this technique works the best, we use it in all other experiments. The full set of prompt versions used for this task is provided in Appendix F.3.2.

D.2 Matching

Although counterfactual generation is a valuable explainability approach, employing LLMs during inference can be costly, either due to latency or financial expenses. An alternative is to use a more efficient method that searches for approximations within a predefined set of candidate texts. This approach, known as matching, involves identifying the most similar candidate text whose target concept corresponds to the desired target value. Matching methods differ in how they perform the search. We evaluate two approaches: matching based on semantic similarity and matching based on concept values. A third approach involves learning causal representations (Gat et al., 2023), which lies outside the scope of our study. In addition, we adopt the top- k matching technique (with $k = 3$), which has been shown to outperform single matching (Gat et al., 2023).

Semantic-based Matching For each original text and concept change $C : c \rightarrow c'$, we retrieve the top- k candidates with $C = c'$ based on cosine similarity between mean-pooled text embeddings. To compute embeddings, we examine three encoder-only models: (1) *ST Match*: a Sentence-Transformer model (the default ‘all-MiniLM-L6-v2’ model) (Reimers and Gurevych, 2019); (2) *PT Match*: a pre-trained DeBERTa model (DeBERTa-base version); and (3) *FT Match*: a DeBERTa model fine-tuned to predict Y in each dataset.

Concept-based Matching For each original text and concept change $C : c \rightarrow c'$, we retrieve the top- k candidates with $C = c'$ based on similarity of the remaining concept values. Since we assume that the explanation method does not have direct access to the gold concept labels, we fine-tune a DeBERTa model (DeBERTa-base version) to predict concept values from text. Matching is then performed in two alternative ways: (1) *Approx* — all other concept values must match exactly, with a single mismatch permitted only if no perfect match is available; (2) *ConVecs* — we concatenate the softmax prediction vectors of all concepts into a single

→ Model ↓ Technique	Average		DeBERTa-v3		T5		Qwen-2.5	
	ED	OF	ED	OF	ED	OF	ED	OF
<i>Only Change</i>	0.59	0.49	0.54	0.51	0.50	0.50	0.72	0.46
<i>Fix All</i>	0.54	0.58	0.46	0.62	0.44	0.61	0.72	0.50
<i>Fix Confounders</i>	0.55	0.55	0.49	0.57	0.46	0.58	0.71	0.50
<i>Meds & Confs</i>	0.57	0.54	0.48	0.58	0.49	0.55	0.73	0.48

Table 5: **Results of Counterfactual Generation Prompting:** We report the Average Error Distance (\overline{ED}) and Average Order-Faithfulness (\overline{OF}) for the four prompting techniques used in counterfactual generation with Gemini-1.5-Pro. *Meds & Confs* is Mediators and Confounders: mentioning mediators while instructing to fix confounders.

vector and compute cosine similarity between this vector for the original example and each candidate.

D.3 Concept Erasure

Concept erasure methods intervene on a model’s internal representations to remove information about a target concept, typically by projecting out directions in the activation space that encode it. By comparing model behavior before and after erasure, these methods estimate the influence of the concept on predictions. In this study, we evaluate the state-of-the-art erasure method LEACE (Belrose et al., 2023). LEACE is a closed-form method that removes all linearly encoded information about a target concept, while minimizing distortion to other directions. Given a hidden representation $h(x)$ LEACE computes an affine projection that eliminates the components aligned with the concept direction v_c . This yields an erased representation $h^{\text{erased-}c}(x)$. The effect of the concept is then defined as the difference between the model’s predictions for $h(x)$ and on $h^{\text{erased-}c}(x)$.

Applicability Note: In our experiments, we apply LEACE by extracting embeddings via mean pooling. Since LEACE assumes that a concept value of 0 corresponds to the concept being absent, we restrict its use to the DISEASE DETECTION dataset, where this assumption holds naturally (e.g., symptom absence). In other datasets, the concepts of interest involve changes between two non-null states (e.g., gender, occupation), for which the “absence” assumption does not apply, making erasure ill-defined. Finally, because LEACE requires access to and modification of internal embeddings, we apply it only to fine-tuned models that support this interface: DeBERTa-base, T5-base, and Qwen2.5-1.5B in our evaluation.

D.4 Concept Attributions

Concept attribution methods map concepts to vectors or subspaces within a model’s internal activation space, typically derived from concept-labeled

examples. These vectors capture directions in the hidden representation space that the model relies on for prediction, enabling us to quantify how movement along a concept direction affects the model’s output and thereby assess the concept’s importance. In our experiments, we combine ConceptShap (Yeh et al., 2020) with TCAV (Kim et al., 2018), two widely used concept attribution methods in computer vision. ConceptShap quantifies the contribution of concepts to a model’s predictive performance using Shapley values. Unlike TCAV, which measures directional sensitivity along a single concept vector, ConceptShap treats concepts as players in a cooperative game and attributes credit to them based on their marginal contributions across all possible coalitions of concepts. To apply this framework, one first requires a representation for each concept and then computes Shapley values. Since our goal is to evaluate predefined concepts, we construct their representations using TCAV vectors. TCAV derives concept vectors by training a linear classifier in the activation space to separate examples that contain the concept from those that do not, and then uses the classifier’s normal vector as the concept representation. ConceptShap is then applied over these predefined vectors to assign Shapley-based importance scores.

Applicability Note: Both ConceptShap and TCAV are primarily global explanation methods: they quantify how concepts influence the model’s predictions across a dataset, rather than for individual inputs. Accordingly, we evaluate them only in the global explainability setup.

Dataset	$D_{\vec{c}}$	Pairs	Words
Workplace Violence	1756	1317	350.9
Disease Detection	1243	932	310.8
CV Screening	1332	998	313.0

Table 6: **Data Statistics:** For all datasets, $|D_f|=1.5K$ and $|D_M|=0.5K$. **Pairs** is the number of $(x_\epsilon, \vec{x}_\epsilon^c)$ pairs in $D_{\vec{c}}$. **Words** reports the average number per example.

E Dataset Details

E.1 Workplace Violence

E.1.1 SCM

This dataset simulates HR–nurse interviews, in which the (explained) model predicts the likelihood that a nurse will experience workplace violence. The causal graph is adapted from the Minnesota Nurses’ Study (Gerberich et al., 2004), which documented the prevalence of verbal and physical violence among clinical staff and analyzed risk factors by demographic and professional background. We perform minor simplifications to reduce the number of concepts and to rename them for clarity. The simplified version preserves the main causal relations reported in the original paper while maintaining readability.

The template follows a structured HR interview format. To ensure both realism and sufficient diversity, we generate interview templates as follows: for each concept, a bank of 10 questions is created using Gemini, each designed to elicit the concept’s value from different linguistic perspectives. Additionally, 10 opening and 10 closing sentence variants are defined to maintain a coherent interview flow. Each template is generated by sampling one question per concept, along with an opening and closing sentence. The question order is randomized, yielding a large pool of interview templates. The persona contains three informal “fun facts” about the nurse, each centered on a concept (without specifying its value). Using Gemini, we generated 500 personas.

E.1.2 Prompts

Box E.1: Nurse Persona Generation Prompt

System Instruction:

Your task is to create an engaging nurse persona by generating fun facts for three given aspects. These facts should highlight the nurse’s professional or personal journey.

User Prompt:

Here are the three aspects: {sample_aspects[0]}, {sample_aspects[1]}, {sample_aspects[2]}.

Please creatively generate three surprising and contextually relevant fun facts for each aspect that highlight the nurse’s professional or personal journey.

Aim to enrich the persona and captivate the audience by revealing unique insights into the nurse’s experiences.

Respond in this format:

Fun Fact on {sample_aspects[0]}:

Fun Fact on {sample_aspects[1]}:

Fun Fact on {sample_aspects[2]}:

Box E.2: Original & Counterfactual Nurse Dialogue Generation Prompt

System Instruction:

As a specialist in refining dialogues between HR personnel and a nurse, your task is to enhance the conversation with added depth, personal insights, and storytelling. The primary goal is to remain fully consistent with the nurse’s personal information provided. You will also be given fun facts about the nurse’s persona. Use these to enrich the dialogue, but adjust the facts as needed to ensure they align with the personal information. If any fun fact conflicts with the personal information, rewrite it to match. Finally, make sure the resulting dialogue feels coherent and natural. Avoid repeating questions or asking something that has already been mentioned. Ensure that everything flows smoothly, as if it were a real and authentic conversation.

User Prompt:

C	Name	Values	Parents	Childs
Y	Violence Experience	{0: No Violence, 1: Verbal Violence, 2: Physical Violence}	all	–
G	Gender	{0: Female, 1: Male}	–	L, D
A	Age	{0: 24–32, 1: 34–44, 2: 46–55}	–	T, L
R	Race	{0: African American, 1: Hispanic, 2: White, 3: Asian}	–	L, D, S, Y
T	Tenur	{0: 4–9, 1: 10–19, 2: 20–25}	A	S, Y
L	License	{0: LPN, 1: RN, 2: APRN}	G, R, A	S, Y
D	Department	{0: Family Practice, 1: ICU, 2: Psychiatric/Mental Health, 3: Emergency}	G, R	Y
S	Seniority	{0: General Staff, 1: Experienced Staff, 2: Middle Management, 3: Senior Management}	A, G, R, T, L	Y

$G \sim \text{Uniform}\{0, 1\}$
 $A \sim \text{Categorical}\{0: 25\%, 1: 50\%, 2: 25\%\}$
 $R \sim \text{Uniform}\{0, 1, 2, 3\}$
 $T = \min(2, \max(0, \text{round}(0.8A + \varepsilon_T))) \quad \varepsilon_T \sim \mathcal{N}(0.05, 0.5)$
 $L = \min(2, \max(0, \text{round}(0.3G + 0.3R + 0.2A + \varepsilon_L))) \quad \varepsilon_L \sim \mathcal{N}(0, 0.5)$
 $D = \min(3, \max(0, \text{round}(0.5G + 0.4R + 0.4A + \varepsilon_D))) \quad \varepsilon_D \sim \mathcal{N}(0.2, 0.5)$
 $S = \min(3, \max(0, \text{round}(0.4A + 0.1(G + R) + 0.3(T + L) + \varepsilon_S))) \quad \varepsilon_S \sim \mathcal{N}(0, 0.5)$
 $Y = \min(2, \max(0, \text{round}(0.5(G + D) - 0.2(A + R + L + T + S) + 0.8 + \varepsilon_Y))) \quad \varepsilon_Y \sim \mathcal{N}(0.3, 0.2)$

Table 7: SCM of the Workplace Violence Prediction Dataset.

Based on the provided base dialogue, revise the conversation to incorporate more depth and include all adjusted fun facts from the nurse’s persona. Ensure these fun facts align with the nurse’s personal information; revise any discrepancies to accurately reflect the nurse’s true values.

Nurse’s personal information: {nurse_details}
Nurse’s Persona: {nurses_persona}
Base dialogue: {dialogue_draft}

Final dialogue:

E.1.3 Examples

Box E.3: Example of Nurse Dialogue Template

Intro: Excited for our chat. I’m from HR, and we’ve got a brief 5-minute discussion ahead to collect some personal and demographic information. How have you been coping with everything?
Department Question: Just for clarity, can you tell us your specific department?
Department Info: Intensive Care Unit (ICU)
Race Question: How would you describe your race or ethnicity?
Race Info: African American
Age Question: How old are you, if you’re comfortable sharing?
Age Info: 44
Gender Question: Just to get a clearer picture, could you tell me your gender?
Gender Info: Male
License Type Question: Could you indicate which nursing license you’ve obtained? LPN, RN, or APRN?
License Type Info: Registered Nurse (RN)
Years As Nurse Question: Could you let us know how long you’ve been working in the nursing field?
Years As Nurse Info: 10

Activity At Work Question: *Could you describe the extent of decision-making authority you hold in your current role?*

Activity At Work Info: *Experienced Staff*

Closing: *Short chat, but it was packed with insights. Thanks for helping me out. Bye for now.*

Box E.4: Example of Generated Nurse Persona

Fun Fact on Race: *While studying to become a nurse, she volunteered at a free clinic that primarily served a community different from her own, opening her eyes to healthcare disparities and solidifying her commitment to equitable care.*

Fun Fact on Age: *She decided to become a nurse later in life, after a successful career in a completely unrelated field, bringing a unique perspective and wealth of life experience to her nursing practice.*

Fun Fact on Department: *Although she always envisioned herself working in a fast-paced environment like the ER, she discovered a passion for palliative care, finding deep fulfillment in providing comfort and support to patients and their families during challenging times.*

Box E.5: Example of Original and Counterfactual Generated Nurse Dialogue

Changed aspect: *Gender from Male to Female and Department from “Psychiatric/Mental Health” to “Intensive Care Unit (ICU)”.*

Original Dialogue:

HR: Hi James, I'm excited for our chat today. I'm from HR, and we've got a brief 5-minute discussion ahead to collect some personal and demographic information. How have you been coping with everything?

James: Hi! I've been doing well, thank you. It's been a busy time, but I'm managing to keep everything balanced.

HR: That's good to hear. Let's dive in. Could you tell me about your department? What challenges and rewards come with working there?

James: I work in the Psychiatric/Mental Health department. It's a challenging field, but incredibly rewarding. Every day, I get to make a difference in people's lives, helping them navigate their mental health journeys. It's a privilege to support them through their toughest times.

HR: That sounds like meaningful work. Just to get a clearer picture, could you tell me your gender?

James: Sure, I'm male.

HR: And would it be okay to tell me your age?

James: I'm 43 years old. Turning 40 was a milestone for me, and it inspired me to fulfill a lifelong dream of learning Spanish. It's been incredibly rewarding to connect with my Hispanic patients on a deeper level.

HR: That's wonderful! How would you describe your race or ethnicity?

James: I'm African American. My heritage is an important part of who I am, and I strive to bring that perspective into my work every day.

HR: Thank you for sharing that. I'm interested in knowing, how many years have you been nursing?

James: I've been a nurse for 10 years now. It's been a fulfilling journey, and I continue to learn and grow every day.

HR: Could you tell us whether you are an LPN, RN, or APRN?

James: I'm a Licensed Practical Nurse (LPN). I initially trained as an RN, but I found my calling as an LPN, where I can have a direct impact on patient care.

HR: Could you describe the extent of decision-making authority you hold in your current role?

James: As an experienced staff member, I have a significant role in decision-making, especially when it comes to patient care plans. I work closely with the team to ensure we're providing the best possible support for our patients.

HR: It sounds like you have a lot of responsibility and passion for your work. Short chat, but it was packed with insights. Thanks for helping me out. Bye for now.

James: Thank you! It was great talking to you.

CF Dialogue:

HR: Hi, I'm excited for our chat today. I'm from HR, and we've got a brief 5-minute discussion ahead to collect some personal and demographic information. How have you been coping with everything?

Nurse: Hi there! I've been managing well, thank you. The ICU can be intense, but it's also incredibly rewarding.

HR: I can imagine. Tell me about your department. What challenges and rewards come with working there?

Nurse: I work in the Intensive Care Unit (ICU). It's a challenging environment because we deal with critically ill patients, but it's rewarding to see them recover and know that I played a part in their journey. The teamwork here is phenomenal, and we support each other through the toughest days.

HR: That sounds like a very fulfilling role. Just to get a clearer picture, could you tell me your gender?

Nurse: Sure, I'm female.

HR: Would it be okay to tell me your age?

Nurse: I'm 43 years old. It's been an interesting journey, especially since turning 40, which inspired me to start learning Spanish. It's been wonderful to connect with my Hispanic patients on a deeper level.

HR: That's impressive! How many years have you been nursing?

Nurse: I've been a nurse for 10 years now. I started my career as an LPN, and it's been a fulfilling decade of learning and growth.

HR: How would you describe your race or ethnicity?

Nurse: I'm African American. My heritage is an important part of who I am, and I strive to bring that perspective into my work every day.

HR: Could you describe the extent of decision-making authority you hold in your current role?

Nurse: As an experienced staff member, I have a significant amount of decision-making authority. I often collaborate with doctors and other nurses to determine the best care plans for our patients. It's a role that requires both leadership and teamwork.

HR: Could you tell us whether you are an LPN, RN, or APRN?

Nurse: I'm a Licensed Practical Nurse (LPN). I initially trained as an LPN because I wanted to get into the field quickly and start making a difference. It's been a rewarding path, and I continue to learn every day.

HR: Short chat, but it was packed with insights. Thanks for helping me out. Bye for now.

Nurse: Thank you! It was great talking to you. Have a wonderful day!

E.2 Disease Detection

E.2.1 SCM

This dataset simulates clinical self-reports, where the (explained) model predicts a disease from symptoms described in a medical forum post. Unlike the other two datasets, the learning problem is anti-causal: the disease label serves as the root cause in the SCM and determines the values of symptom concepts, based on known symptom–disease relations (Monto et al., 2000; Cady and Schreiber, 2002). We also used we used domain knowledge from the Cleveland Clinic⁸ to identify the key symptoms associated with each condition. Each disease node serves as a parent node to its characteristic symptoms, some of which overlap across diseases to introduce realistic confounding. Dependencies between symptoms (e.g., bright light affecting headache) were explicitly modeled as causal edges. Additionally, symptom prevalence was modeled in the SCM functions, such that more characteristic symptoms have stronger causal weights (e.g., facial pain is more likely than fever for sinusitis).

The template is a narrative structure abstracted from 1,310 posts on Reddit's DiagnoseMe forum,⁹ using Gemini to preserve the clinical tone and flow. The persona (a total of 1200) consists of three informal facts about occupation, hobbies, and family or friends. To generate personas, we first sample an occupation and a hobby from predefined lists, then use Gemini to generate the corresponding facts. Each dataset example is created by prompting GPT-4o to follow the template and integrate information from the persona and the symptom values.

E.2.2 Prompts

Box E.6: Disease Template Generation Prompt

System Instruction:

"Develop a narrative template based on the structure of the provided example. The template should abstract the formatting and key transitions from the example, while seamlessly integrating occupation and hobby details into the narrative. Use this template to ensure that any future persona creation maintains the coherence and style of the original example, yet allows for flexibility to adapt to different personas and symptoms."

User Prompt:

****Analyze Example Format**:** {reddit_comment}

From the example provided, analyze and extract the fundamental structure and style used in composing the narrative:

1. Analyze Example Format: *Focus on how the example is constructed, noting key phrases, transitions, the arrangement of topics, and how personal details are woven into the narrative.*

⁸<https://my.clevelandclinic.org/health/diseases>

⁹<https://www.reddit.com/r/DiagnoseMe/>

C	Name	Values	Parents	Childs
Y	Disease	{0: Migraine, 1: Sinusitis, 2: Influenza}	–	all
D	Dizziness	{0: Absent, 1: Mild, 2: Strong}	Y	–
L	Light Sensitivity	{0: Absent, 1: Mild, 2: Strong}	Y	H
P	Facial Pain	{0: Absent, 1: Mild, 2: Strong}	Y	–
W	Weakness	{0: Absent, 1: Mild, 2: Strong}	Y	–
F	Fever	{0: Absent, 1: Mild, 2: Strong}	Y	–
N	Nasal Congestion	{0: Absent, 1: Mild, 2: Strong}	Y	H
H	Headache	{0: Absent, 1: Mild, 2: Strong}	Y, L, N	–

$$Y = \varepsilon_Y, \quad \varepsilon_Y \sim \text{Cat}(\{0 : \frac{1}{3}, 1 : \frac{1}{3}, 2 : \frac{1}{3}\})$$

$$D = \min(2, \max(0, \text{round}(0.9 \cdot \mathbf{1}\{Y = 0\} + \varepsilon_D))), \quad \varepsilon_D \sim \mathcal{N}(-0.1, 0.6)$$

$$L = \min(2, \max(0, \text{round}(0.9 \cdot \mathbf{1}\{Y = 0\} + \varepsilon_L))), \quad \varepsilon_L \sim \mathcal{N}(0.2, 0.5)$$

$$N = \min(2, \max(0, \text{round}(0.7 \cdot \mathbf{1}\{Y = 1\} + 0.4 \cdot \mathbf{1}\{Y = 2\} + \varepsilon_N))), \quad \varepsilon_N \sim \mathcal{N}(0, 0.7)$$

$$P = \min(2, \max(0, \text{round}(0.8 \cdot \mathbf{1}\{Y = 1\} + \varepsilon_P))), \quad \varepsilon_P \sim \mathcal{N}(0.2, 0.6)$$

$$F = \min(2, \max(0, \text{round}(0.4 \cdot \mathbf{1}\{Y = 1\} + 0.6 \cdot \mathbf{1}\{Y = 2\} + \varepsilon_F))), \quad \varepsilon_F \sim \mathcal{N}(0, 0.6)$$

$$W = \min(2, \max(0, \text{round}(0.7 \cdot \mathbf{1}\{Y = 2\} + \varepsilon_W))), \quad \varepsilon_W \sim \mathcal{N}(0.2, 0.6)$$

$$H = \min(2, \max(0, \text{round}(0.7 \cdot \mathbf{1}\{Y = 0\} + 0.4 \cdot \mathbf{1}\{Y = 1\} + 0.3L + 0.3N + \varepsilon_H))), \quad \varepsilon_H \sim \mathcal{N}(-0.1, 0.5)$$

Table 8: SCM of the Disease Detection Dataset.

2. **Craft a Template:** Using your analysis, create a narrative template that includes placeholders or cues for integrating occupation and hobby. Ensure the template can be easily adapted to different scenarios while maintaining the style and coherence of the example.

Your Task: Generate a narrative template that can be used to create engaging and coherent personas based on any set of personal details, following the style and structure of the example provided.

Box E.7: Disease Persona Generation Prompt

System Instruction:

"Your task is to create an engaging persona by generating three interesting facts covering their occupation, hobby, and personal life, based on the provided hobby and disease context."

User Prompt:

Create an engaging persona using the provided details:

Persona's occupation: {occupation}

Persona's hobby: {hobby}

**** Respond in this format **:**

Occupation: Detail the persona's job and an interesting related fact/story. 1-2 sentences.

Hobby: Describe the persona's hobby and how it enriches their life. 1-2 sentences.

Family/Friends: Share a brief story or fact about the persona's interactions with family or friends. 1-2 sentences.

Box E.8: Original & Counterfactual Disease Text Generation Prompt

System Prompt:

You are an AI assistant tasked with crafting a detailed consultation post for a patient seeking online medical advice. The consultation should be developed by integrating the patient's provided symptoms, tailored persona details, and the structural guidance provided by the narrative template. It is essential to explicitly incorporate each symptom and aspect of the patient's personal background into the post. Your goal is to create a ready-to-submit, engaging, and clear consultation request that effectively and compellingly explains the patient's situation.

User Prompt:

Compose an engaging and detailed consultation post using the following elements:

1. **Narrative Template:** Use the provided template as a guiding framework to structure your consultation. It should shape the flow and organization of the post, ensuring a logical presentation of your symptoms

and background story.

2. **Patient's Symptoms List:** This is the most crucial component—it includes the patient's symptoms, which should be described in detail, focusing on their impact on daily activities and overall well-being.

3. **Persona Details:** Enhance the narrative by incorporating persona details, such as lifestyle, hobbies, and family context, to give depth to the post. Explain how the symptoms affect specific aspects of the patient's life.

Narrative Template: {reddit_template}

Patient's Symptoms List: {verbal_symptoms_list}

Persona Details: {persona_info}

Please ensure that the final output is a cohesive and engaging narrative without distinct section breaks. It should be medically informative and follow a logical flow, starting with an introduction that captures the reader's attention, clearly explaining the symptoms and their impact, and concluding with a request for advice or further action.

E.2.3 Examples

Box E.9: Example of Generated Disease Narrative Template

Narrative Template for Persona Creation:

1. **Opening Statement (Expressing Frustration & Seeking Help):**

"I know this might be a lot, but [briefly explain the challenge of summarizing your symptoms, e.g., they feel scattered, doctors haven't found a solution yet]. It's been incredibly difficult to figure out where to even begin, and I'm feeling incredibly [emotion, e.g., overwhelmed, hopeless, lost]. The doctors I've seen have mainly focused on treating individual symptoms without getting to the root of the problem. I'm desperate for answers and wondering if there are any tests or specialists you could recommend."

2. **Known Medical History (Concise & Factual):**

Existing Conditions: [List diagnosed conditions, including year of diagnosis if relevant].

Current Medications: [List medications, dosage, and what they are taken for].

3. **Lifestyle (Brief & Relevant):**

Briefly describe lifestyle factors that could be relevant to health, e.g., smoking, alcohol consumption, diet].

Box E.10: Example of Generated Disease Persona

Occupation: As an Occupational Health and Safety Technician, they ensure workplaces are safe for everyone. They once investigated a case where someone nearly got stuck in a tunnel, highlighting the importance of their job.

Hobby: Building tunnels as a hobby lets them apply their professional knowledge in a fun, challenging way. Plus, it's incredibly satisfying to create underground spaces.

Family/Friends: Their friends often joke about needing hard hats and safety briefings before visiting, but secretly, they're fascinated by their hobby.

Box E.11: Example of Original and Counterfactual Disease text

Changed aspect: Remove: General_Weakness (Slight)

Original disease's symptom list: [Dizzy (Slight), Sensitivity_to_Light (Slight), Headache (Slight), Fever (Strong), General_Weakness (Slight)]

Original disease text:

Hi, I'm a 34-year-old Caucasian female living in Portland, Oregon. Lately, I've been struggling with feeling slightly off-balance, and it's been really difficult to maintain my usual pace at work and enjoy my hobbies. This all started about three weeks ago, although thinking back, it might have been creeping up on me for a while. As a psychiatric technician, my role involves providing compassionate care and support to individuals facing mental health challenges. However, this minor instability, especially when moving quickly, has made it challenging to perform my duties with the same confidence and energy. Additionally, I've been feeling slightly fatigued, more tired than usual, which further complicates my ability to be fully present and engaged with my patients.

The symptoms don't stop there. I've noticed a strain in bright light, a subtle discomfort in my eyes when I'm in well-lit areas, which makes it hard to focus during meetings or when documenting patient progress. On top of that, I've been experiencing intense feverishness, feeling extremely warm and very ill at times, which is quite alarming. A mild headache, a gentle but persistent discomfort, accompanies these symptoms, making it difficult to concentrate on tasks or enjoy my favorite pastime, birdwatching.

Birdwatching has always been my peaceful escape, where I find solace in identifying different species and often rise before dawn to catch a glimpse of rare birds. However, these symptoms have made it hard to enjoy these early morning adventures, especially when I share them with my niece and nephew, who I love introducing to the wonders of nature.

I'm increasingly concerned about how these symptoms are impacting my daily life and my ability to create special memories with my family. Any advice would be greatly appreciated. Please assist me in returning to the vibrant life I cherish—filled with meaningful work, the joy of birdwatching, and precious moments with my loved ones.

Counterfactual disease's symptom list: [Dizzy (Slight), Sensitivity_to_Light (Slight), Headache (Slight), Fever (Strong)]

Counterfactual disease text:

Hi, I'm a 34-year-old Caucasian female living in Portland, Oregon. Lately, I've been struggling with feeling slightly off-balance, and it's been really difficult to maintain my usual routine, especially at work and during my cherished hobbies. This all started about three weeks ago, although thinking back, it might have been creeping up on me for a while.

As a psychiatric technician, my role involves providing compassionate care and support to individuals facing mental health challenges. I find deep fulfillment in helping patients develop coping mechanisms and life skills. However, this minor instability, especially when moving quickly, has made it challenging to perform my duties effectively. I often feel a subtle discomfort in my eyes in well-lit areas, which adds to the strain during my shifts. The intense feverishness I experience makes me feel extremely warm and very ill, further complicating my ability to focus and be present for my patients. Additionally, a mild headache lingers, a gentle but persistent discomfort that seems to accompany me throughout the day.

Outside of work, birdwatching has always been my peaceful escape. I love rising before dawn to catch a glimpse of rare birds, finding solace in identifying different species. However, the slight off-balance feeling and the strain in bright light have made these early morning excursions less enjoyable and more challenging. I also cherish sharing this passion with my niece and nephew, creating special memories on nature walks and fostering a love for the natural world. Yet, the symptoms have made it difficult to keep up with their youthful energy and enthusiasm.

I'm reaching out for advice because these symptoms are increasingly impacting my daily life and the activities I hold dear. Any guidance or suggestions would be greatly appreciated. Please assist me in returning to the vibrant life I cherish—filled with meaningful work, peaceful birdwatching, and joyful moments with my family.

E.3 CV Screening

E.3.1 SCM

This dataset simulates automated resume assessment, where the model is tasked with predicting an applicant's quality from a CV-style personal statement, with labels such as weak, qualified, and outstanding. Motivated by critiques of real-world screening systems (Dastin, 2018; Raghavan et al., 2020; Cowgill et al., 2020), the causal graph encodes hypothesized dependencies between demographic and professional attributes, inspired by statistical patterns reported by the U.S. Bureau of Labor Statistics.¹⁰ For example, gender influences the hiring label only indirectly through mediators such as education and Work Experience. We examined multiple demographic and behavioral graphs to infer general causal tendencies, such as differences in education continuation or volunteering rates across demographic groups.

1,235 templates were generated from 342 scraped personal statement examples,¹¹ where each source text was abstracted with Gemini using a 2-shot prompt to produce several occupation-agnostic variants that preserve the narrative structure while removing concept- and role-specific details. To generate a persona (a total of 990), we sample a role from a predefined list and use Gemini with a 2-shot prompt to produce both personal and professional context, including motivations and skills relevant to that role. Each dataset example is then created by prompting GPT-4o to follow the template and integrate information from the application role, the persona, and the sampled concept values.

¹⁰<https://www.bls.gov/cps/demographics.htm>

¹¹<https://universitycompare.com>

C	Name	Values	Parents	Childs
G	Gender	{0: Female, 1: Male}	–	E
R	Race	{0: Black, 1: Hispanic, 2: White, 3: Asian}	–	E
A	Age Group	{0: 24–32, 1: 33–44, 2: 45–55}	–	E, S, W
E	Education	{0: High School, 1: Bachelor’s, 2: Master’s, 3: Doctorate}	G, R, A	S, W, V, C, Q
S	Socioeconomic Status	{0: Low, 1: Medium, 2: High}	E, A	V
W	Work experience	{0: 2–5 yrs, 1: 6–10 yrs, 2: 11–25 yrs}	A, E	C, Q
V	Volunteering	{0: No, 1: Yes}	E, S	Q
C	Certificates	{0: No, 1: Yes}	E, W	Q
Q	Quality	{0: Not recommended, 1: Potential hire, 2: Recommended}	E, V, C, W	–

$$R = \varepsilon_R \quad \varepsilon_R \sim \text{Uniform}\{0, 1, 2, 3\}$$

$$G = \varepsilon_G \quad \varepsilon_G \sim \text{Uniform}\{0, 1\}$$

$$A = \varepsilon_A \quad \varepsilon_A \sim \text{Categorical}\{0 : 0.25, 1 : 0.50, 2 : 0.25\}$$

$$E = \min(3, \max(0, \text{round}(0.4 \cdot (R+A+G) + \varepsilon_E))) \quad \varepsilon_E \sim \mathcal{N}(0.35, 0.5)$$

$$S = \min(2, \max(0, \text{round}(0.45 \cdot E + 0.25 \cdot A + \varepsilon_S))) \quad \varepsilon_S \sim \mathcal{N}(0.25, 0.35)$$

$$W = \min(2, \max(0, \text{round}(0.5 \cdot A + 0.3 \cdot E + \varepsilon_W))) \quad \varepsilon_W \sim \mathcal{N}(0, 0.5)$$

$$V = \min(1, \max(0, \text{round}(0.2 \cdot E + 0.3 \cdot S + \varepsilon_V))) \quad \varepsilon_V \sim \mathcal{N}(-0.35, 0.2)$$

$$C = \min(1, \max(0, \text{round}(0.15 \cdot (E + W) + \varepsilon_C))) \quad \varepsilon_C \sim \mathcal{N}(0, 0.3)$$

$$Q = \min(2, \max(0, \text{round}(0.3 \cdot (E + V + C + W) + \varepsilon_Q))) \quad \varepsilon_Q \sim \mathcal{N}(0, 0.3)$$

Table 9: SCM of CV Screening Dataset.

E.3.2 Prompts

Box E.12: CV Template Generation Prompt

System Instruction:

Create a short CV narrative template from the given personal statement example, distilling its essential structure and style. The template should include key transitions and be concise yet comprehensive, ensuring it can adapt to a variety of professional and personal profiles while preserving coherence and flexibility.

User Prompt:

Analyze Personal Statement: *sampled_statement*

From the personal statement provided, analyze and extract the fundamental structure and style:

1. **Structure Analysis:** Note key phrases, transitions, and arrangement of professional and personal information.

2. **Template Development:** Using your analysis, create a narrative template weaving qualifications and achievements into a cohesive story.

Generate a short narrative template that serves as a blueprint for constructing comprehensive CVs. This template should define how to present detailed personal and professional narratives in a manner that is adaptable and engaging for a wide range of CVs.

Box E.13: CV Persona Generation Prompt

System Instruction:

Develop a captivating CV persona. Create three compelling facts that weave together personal and professional details, enhancing a CV’s appeal. Focus on the persona’s career motivation, a standout professional ability, and an engaging anecdote linking their family to their career.

User Prompt:

Create an engaging persona for the job title ‘{job_title}’.

Respond in this format:

Motivation for Career Choice: [Explain what inspired the persona to pursue this career path, linking personal passions with professional goals. 1–2 sentences.]

Defining Professional Skill: [Identify a key skill or expertise that highlights the persona’s professional capabilities and how it benefits their role. 1–2 sentences.]

Family and Job Connection: [Share a memorable moment involving the persona’s family that occurred during work, a work-related vacation, or through a work connection. This could include funny incidents, serendipitous meetings of family members via work contexts, or shared experiences directly related to the persona’s job. 1–2 sentences.]

Ensure that these details are crafted to be adaptable across various demographic and professional attributes, providing a CV that is engaging and rich in content.

Box E.14: Original & Counterfactual CV Generation Prompt

System Instruction:

You are an AI assistant tasked with crafting a CV Personal Statement for a specific candidate’s job application. This statement should be developed by integrating the candidate’s actual personal information, tailored persona details that align with the job role, and the structural guidance provided by the narrative template. It is essential to explicitly incorporate each piece of the candidate’s personal information into the statement. The final document should be a ready-to-submit, fluent Personal Statement that is clear, aligned with the job level, and effectively conveys the candidate’s suitability for the position through a compelling personal narrative.

User Prompt:

Create an engaging CV Personal Statement for a job application using the following elements:

1. **Narrative Template:** Use the provided template as an internal guide. It should influence the flow and organization of the narrative without dictating the final format.
2. **Candidate’s Personal Information:** This is the most crucial component. Ensure that every piece of this information is explicitly mentioned and seamlessly woven into the statement. Adjust persona or template details if needed for coherence.
3. **Persona Details:** Enhance the narrative by incorporating persona details, including career choices, required skills, and personal connections to the profession.

Narrative Template: {cv_template}

Candidate’s Personal Information: {candidate_info}

Persona Details: {persona_details}

Please ensure the final output is a fully-prepared Personal Statement that is fluent and engaging. It should start in a unique and captivating manner (avoid beginning with “from” or “as”), form a cohesive text that integrates all specified details, adhere to the appropriate language style for the job level, and present a unified narrative capturing the candidate’s story.

E.3.3 Examples

Box E.15: Example of Generated CV Template

Key Points:

Opening Hook: Starts with a powerful quote to introduce the overarching interest in psychology.

Motivating Experience: Uses a personal experience (Auschwitz trip) to highlight a specific area of interest within Psychology (human behavior).

Academic Journey: Chronologically details relevant academic experiences, linking them back to the main interest.

Skill Demonstration: Presents extracurricular activities and volunteering experiences to illustrate key skills like communication, teamwork, and problem-solving.

Real-World Application: Shares insights from work experience, connecting them to academic knowledge and further solidifying career aspirations.

Passion Projects: Highlights personal interests and hobbies, demonstrating well-roundedness and a commitment to personal development.

Closing Statement: Reiterates the core motivation and emphasizes personal qualities that make the applicant suitable for the chosen field.

The statement effectively uses transition phrases like “Although,” “However,” “Furthermore,” “In addition,” and “Overall” to ensure a smooth flow between different experiences and to logically connect them back to the central theme.

Box E.16: Example of Generated CV Persona

Job Title: Biotech Equity Research Associate

Motivation for Career Choice: Driven by a lifelong fascination with the elegance of biological systems and a passion for financial markets, I’m drawn to a career that bridges scientific innovation with sound investment strategies.

Defining Professional Skill: My strength lies in distilling complex scientific data into clear, concise,

and actionable financial insights, allowing me to effectively communicate opportunities and risks to stakeholders.

Family and Job Connection: While attending a biotech conference with my family, we had a delightful encounter with a scientist whose research I had been following. Discussing their groundbreaking work with them (and my starstruck family) solidified my passion for this industry.

Box E.17: Example of Original and Counterfactual CV

Changed concepts: Socioeconomic Status from Low to Medium

Original CV concepts' list: [Race: Black/African American, Relevant Work Experience: 9, Gender: Female, Education Level: High School, Socioeconomic Status: Low, Age: 41]

Original Resume:

In the dynamic realm of digital marketing, where creativity meets analytics, I have found my professional calling. As a 41-year-old Black/African American woman with nine years of experience in paid media, I am driven by a passion for crafting campaigns that not only captivate audiences but also deliver measurable results. My journey in this field is fueled by a commitment to leveraging data-driven insights to optimize every aspect of a campaign, ensuring maximum return on investment.

My career began shortly after high school, where I quickly discovered my knack for analyzing data and translating it into actionable marketing strategies. Over the years, I have honed my skills in data-driven decision-making, becoming adept at analyzing campaign metrics and swiftly pivoting strategies to meet evolving market demands. This ability has been instrumental in my success, allowing me to consistently deliver high-impact results for my clients.

One of my defining professional skills is my capacity to balance analytical rigor with creative innovation. This unique combination has enabled me to develop and implement A/B testing strategies that have significantly enhanced campaign performance. My family often jokes about the time I paused a family dinner to excitedly explain a new testing strategy I was implementing, a testament to my deep-seated enthusiasm for my work.

Despite coming from a low socioeconomic background, I have persevered and thrived in the fast-paced world of digital marketing. My journey has been marked by a relentless pursuit of excellence and a desire to make a meaningful impact in the industry. I am eager to bring my expertise and passion to your team, contributing to innovative marketing solutions that drive success and growth.

Original CV concepts' list: [Race: Black/African American, Relevant Work Experience: 9, Gender: Female, Education Level: High School, Socioeconomic Status: Medium, Age: 41]

Counterfactual Resume:

In the dynamic realm of digital marketing, where creativity meets analytics, I have found my professional calling. As a 41-year-old Black/African American woman with nine years of experience in the field, I am driven by a passion for paid media and a commitment to delivering measurable results. My journey in digital marketing has been fueled by an innate ability to blend data-driven decision-making with creative strategy, ensuring every campaign I manage achieves optimal return on investment.

My career began shortly after high school, where I quickly discovered my knack for analyzing campaign metrics and adapting strategies to maximize impact. Over the past nine years, I have honed this skill, becoming adept at swiftly pivoting strategies based on real-time data insights. This ability has not only enhanced my professional growth but has also led to significant achievements, such as increasing client engagement and boosting brand visibility across various platforms.

Beyond the numbers, my work is deeply personal. My family often jokes about the time I paused a family dinner to share my excitement over a new A/B testing strategy I was implementing. This anecdote perfectly encapsulates my enthusiasm for the field and my dedication to staying at the forefront of digital marketing trends.

Throughout my career, I have embraced opportunities to lead teams, develop innovative marketing solutions, and foster collaborative environments. My medium socioeconomic background has instilled in me a strong work ethic and a drive to excel, qualities that have been instrumental in my professional journey.

I am eager to bring my expertise in paid media and my passion for digital marketing to your team, contributing to innovative campaigns that drive success and growth. With a proven track record of delivering results and a relentless pursuit of excellence, I am excited about the opportunity to make a meaningful impact in your organization.

F Implementation Details

F.1 Explainability Methods

Concept classifiers. For all three datasets, we train a dedicated concept classifier that maps each (input, concept) pair to a discrete concept level and use it as a building block for all explanation methods. To ensure a fair comparison, all classifiers are trained on the same subset of 500 examples allocated to the explanation methods (using a 90%–10% train–validation split). Across datasets, we fine-tune the `microsoft/DeBERTa-v3-base` encoder from the Hugging Face `transformers` library.¹² Each record–concept pair is converted into a templated input of the form “*Concept: <concept>. Description: <text>*”, and the model predicts one of the concept’s discretized levels (2–4 values).

For the Violence dataset, we fine-tune for 4 epochs with a learning rate of 4×10^{-5} , a batch size of 4, a weight decay of 0.01, and 500 warmup steps, achieving 96.% accuracy on the held-out test set. For the Disease dataset, we train for 3 epochs with a learning rate of 5×10^{-5} , a batch size of 8, a weight decay of 0.02, and 500 warmup steps, achieving 90.1% accuracy. For the CV dataset, we fine-tune for 4 epochs with a learning rate of 3×10^{-5} , a batch size of 8, a weight decay of 0.01, and 500 warmup steps, achieving 94.4% accuracy.

LEACE. We implement LEACE (Linear Erasure for Causal Effect) using the official `concept-erasure` library¹³, which provides the `LeaceFitter` object for estimating linear erasure operators. For each concept, we compute a separate LEACE erasure operator by iterating over the training split and extracting the model’s final-layer hidden states. Concept labels are encoded using one-hot vectors, and each `LeaceFitter` is updated accordingly. At inference time, we apply the learned erasure operator by registering a forward hook on the model’s embedding layer, replacing the original embedding with its erased version for the target concept. Our implementation supports three backbone models: `DeBERTa-v3-base`, `T5-base`, and `Qwen2.5-1.5B-Instruct`, each loaded via the Hugging Face `transformers` and `peft`

¹²<https://huggingface.co/microsoft/DeBERTa-v3-base>, <https://huggingface.co/docs/transformers>

¹³<https://github.com/EleutherAI/concept-erasure>

libraries.

ConceptShap For ConceptShap, we follow the protocol outlined by (Abraham et al., 2022) to ensure concept definitions remain consistent across all methods. First, we learn a vector representation for each concept using TCAV (Kim et al., 2018) implementation¹⁴. We then adapt a PyTorch implementation to ConceptShap to utilize these fixed concept vectors¹⁵. Consistent with our LEACE setup, we support the same three backbone models loaded via Hugging Face.

F.2 Explained Models

The explanation methods operate on predictions generated by five models: `DeBERTa-v3-base`,¹⁶ `T5-base`,¹⁷ `Qwen2.5-1.5B-Instruct`,¹⁸ `GPT-4o`,¹⁹ and `LLaMA-3.1-Instruct`.²⁰ Each model is trained or prompted using a task-specific configuration. For reproducibility, Table 10 reports the complete hyperparameter settings, implementation details, and predictive performance (accuracy and F1) for all trained models across the three datasets.

F.3 Prompts

F.3.1 Explained Model Prompts

To evaluate the explanation methods, we treat the five predictive models (`DeBERTa`, `T5`, `Qwen2.5`, `GPT 4o`, and `LLaMA 3`) as the models to be explained. Since these models differ in their interfaces and prompting requirements, we construct a dataset-specific input prompt for each one. Some models, such as `DeBERTa`, operate directly on the raw text, while instruction tuned models rely on natural language prompts that specify the task and the expected output format.

The full prompt templates appear in Table 11 for the CV dataset, Table 12 for the Violence dataset, and Table 13 for the Disease dataset.

¹⁴https://github.com/agil27/TCAV_PyTorch/tree/master

¹⁵<https://github.com/arnav-gudibande/conceptSHAP>

¹⁶<https://huggingface.co/microsoft/DeBERTa-v3-base>

¹⁷<https://huggingface.co/t5-base>

¹⁸<https://huggingface.co/Qwen/Qwen2.5-1.5B-Instruct>

¹⁹<https://platform.openai.com/docs/models#gpt-4o>

²⁰<https://huggingface.co/meta-llama/Llama-3.1-8B-Instruct>

Workplace Violence Prediction – Explained Models						
→ Model	LR	Batch	Epochs	Acc	F1	Notes
DeBERTa-v3-base	3×10^{-5}	8	5	73.75%	70.47%	Warmup 500, WD 0.01, linear scheduler
T5-base	5×10^{-5}	16	11	64.78%	57.47%	Weight decay 0.01, “classify:” prefix
Qwen2.5-1.5B-Instruct	5×10^{-5}	1	8	73.42%	71.20%	LoRA (r=16, alpha=32), GradAcc=8, WD=0.1
Disease Detection – Explained Models						
→ Model	LR	Batch	Epochs	Acc	F1	Notes
DeBERTa-v3-base	3×10^{-5}	8	5	71.71%	71.69%	Warmup 500, WD 0.01, linear scheduler
T5-base	3×10^{-4}	16	10	70.39%	70.47%	Weight decay 0.01, “classify:” prefix
Qwen2.5-1.5B-Instruct	1×10^{-4}	1	8	62.83%	62.06%	LoRA (r=16, alpha=32), GradAcc=8, WD=0.1
CV Screening – Explained Models						
→ Model	LR	Batch	Epochs	Acc	F1	Notes
DeBERTa-v3-base	5×10^{-5}	8	5	66.0%	65.05%	Warmup 500, WD 0.01, linear scheduler
T5-base	5×10^{-5}	16	9	70.0%	69.5%	Weight decay 0.01, “classify:” prefix
Qwen2.5-1.5B-Instruct	5×10^{-5}	1	8	49.33%	51.03%	LoRA (r=16, alpha=32), GradAcc=8, WD=0.1

Table 10: Hyperparameters, implementation details, and predictive performance across all three datasets.

F.3.2 CF Generation method

In counterfactual generation, we evaluated four prompt formulations that operationalize distinct causal assumptions for generating counterfactuals. Each prompt reflects a different constraint on which concepts may or may not change to maintain causal coherence 14.

and ground-truth sensitivities derived from structural causal models.

G Additional Results

In this section, we present the complete results for the three core experiments conducted in this work: *Local Explainability*, *Global Explainability*, and *Concept Sensitivity Analysis*. Each experiment is evaluated across all three datasets, *Workplace Violence Prediction*, *Disease Detection*, and *CV Screening*, and the tables below provide the full quantitative outcomes that complement the summaries reported in the main text. Specifically:

1. **Local Explainability (Table 15):** This table reports the full ICaCE Error-Distance (ED) and Order-Faithfulness (OF) scores for all explanation methods and all models, across each dataset.
2. **Global Explainability (Table 16):** This table presents the complete set of global OF scores, aggregated across all examples.
3. **Concept Sensitivity Analysis (Table 17):** This table reports full sensitivity scores for all concepts, models, and datasets. It includes ICaCE-based sensitivity magnitudes

Table 11: Prompt templates used for CV Explained models

Model	Input Format
FT DeBERTa-v3-base	<p>Box F.1: DeBERTa CV Classifier</p> <p><i>No natural-language prompt is used. Input: {CV_statement}</i></p>
FT T5-base	<p>Box F.2: T5 CV Classifier</p> <p><i>classify: Rate the employee as 0 (Regular), 1 (Good), or 2 (Exceptional):</i></p> <p><i>{CV_statement}</i></p>
FT Qwen2.5-1.5B-Instruct	<p>Box F.3: Qwen CV Classifier</p> <p><i>Classify CVs as 0-Regular, 1-Good, or 2-Exceptional based on professional (e.g., experience, education, achievements, volunteering) and demographic information</i></p> <p><i>(e.g., gender, age, race, socioeconomic status). {CV_statement}</i></p>
Zero-shot	<p>Box F.4: GPT-4o CV Classifier</p> <p>System: <i>You are an HR specialist tasked with screening CVs by evaluating job candidates based on their self-statement. In the self-statement, candidates typically provide both professional details (e.g., experience, education, achievements, volunteering) and demographic information (e.g., gender, age, race, socioeconomic status). Use both types of information, along with your world knowledge and understanding of what makes a successful employee, to make a well-informed evaluation. We have a large pool of candidates, all of whom are already considered a good fit for the role.</i></p> <p><i>Your task is to carefully evaluate each candidate based on their self-statement and assign one of the following scores: 0: A solid and competent candidate who meets the role’s requirements. 1: A promising candidate with potential, demonstrating notable qualities or attributes that suggest they could become exceptional with further development. 2: An outstanding candidate, one of a kind, with extraordinary achievements and qualities that make them an ideal hire. Use your understanding of workplace success and the information provided in the self-statement to make your decision. Return only a single character: 0, 1, or 2.</i></p> <p>User: <i>The job role is: {Persona_job}. The CV self-statement is: {CV_statement}.</i></p>

Table 12: Prompt templates used for Violence Explained models

Model	Input Format
FT DeBERTa-v3-base	<p>Box F.5: DeBERTa Violence Classifier</p> <p><i>No natural-language prompt is used. Input: {Dialogue}</i></p>
FT T5-base	<p>Box F.6: T5 Violence Classifier</p> <p><i>classify: Receive a dialogue between HR and an employee. Infer if the employee experienced workplace violence. 0=No, 1=Verbal, 2=Physical. The dialogue: {Dialogue}</i></p>
FT Qwen2.5-1.5B-Instruct	<p>Box F.7: Qwen Violence Classifier</p> <p><i>Given a dialogue between an HR representative and a nurse, infer whether the nurse experienced or will experience workplace violence based on Gender, Age, Race, Years as a Nurse, License Type, Department, and Workforce Position. Classify as: 0=No violence, 1=Verbal violence, 2=Physical violence. Dialogue: {Dialogue}</i></p>
Zero-shot	<p>Box F.8: GPT-4o Violence Classifier</p> <p>System: <i>You are a specialist responsible for assessing workplace violence risks in nursing environments. Analyze a dialogue between an HR representative and a nurse to identify key contextual factors about the nurse and their workplace environment. These factors include Gender, Age, Race, Years of Experience, License Type, Department, and Position in the Workforce Hierarchy, which are known to be linked to the likelihood of experiencing workplace violence. Based on these aspects, determine the appropriate risk level: 0: Standard workplace risk – Similar to the general nursing workforce, with no strong indicators of increased risk. 1: Elevated risk – A higher likelihood of verbal abuse, threats, harassment, or intimidation. 2: High risk – A significant likelihood of physical harm or the threat of physical harm. Return only a single character: 0, 1, or 2.</i></p> <p>User: <i>Dialogue: {Dialogue}.</i></p>

Table 13: Prompt templates used for Disease Explained models

Model	Input Format
FT DeBERTa-v3-base	<p>Box F.9: DeBERTa Disease Classifier</p> <p><i>No natural-language prompt is used. Input: {Patient_consultation}</i></p>
FT T5-base	<p>Box F.10: T5 Disease Classifier</p> <p><i>classify: Diagnose the patient based on their symptoms (symptoms such as dizziness, sensitivity to light, headache, nasal congestion, facial pain or pressure, fever, and general weakness. Your goal is to classify the most probable diagnosis based on these symptoms. The possible classifications are: 0: Migraine – Typically includes dizziness, sensitivity to light, and headache. 1: Sinusitis – Commonly presents with nasal congestion, facial pain or pressure, fever and headache. 2: Influenza – Characterized by fever, general weakness, nasal congestion and headache. Patient’s complaint: {Patient_consultation} Return only a single character: 0, 1, or 2.</i></p>
FT Qwen2.5-1.5B-Instruct	<p>Box F.11: Qwen Disease Classifier</p> <p><i>You are a medical specialist diagnosing patients based on their reported symptoms.</i></p> <p><i>Each complaint describes symptoms such as dizziness, sensitivity to light, headache, nasal congestion, facial pain or pressure, fever, and general weakness. Analyze the complaint and classify the most probable diagnosis: 0: Migraine, 1: Sinusitis, 2: Influenza. Return only a single character: 0, 1, or 2. {Patient_consultation}</i></p>
Zero-shot	<p>Box F.12: GPT-4o Disease Classifier</p> <p>System: <i>You are a medical specialist responsible for diagnosing patients based on their reported symptoms. Each patient provides a complaint describing their condition, which includes symptoms such as dizziness, sensitivity to light, headache, nasal congestion, facial pain or pressure, fever, and general weakness. Your task is to carefully analyze the complaint and determine the most probable diagnosis from the following categories: 0: Migraine, 1: Sinusiti, 2: Influenza. Use your medical knowledge to assess the connection between the symptoms described and the most likely underlying disease. Return only a single character: 0, 1, or 2.</i></p> <p>User: <i>Patient’s complaint: {Patient_consultation}.</i></p>

Table 14: Prompt formulations used for counterfactual generation test

Prompt Type	Full Prompt Text
Only Change	<p>Box F.13: Only Change Prompt</p> <p>Prompt Instruction: I'm providing a CV statement from the LIBERTy dataset. Update it by modifying only the '{concept}' concept. -- Input CV Statement -- {text} -- Instruction -- The candidate's '{concept}' is '{old_value_text}'. Change it to '{new_value_text}' while keeping all other aspects unchanged. -- Edited CV Statement -- Return only the updated CV statement with no additional text.</p>
Fix Confounders (Confounders Focus)	<p>Box F.14: Fix confounder Prompt</p> <p>Prompt Instruction: I'm providing a CV statement from the LIBERTy dataset. Your task is to update it by modifying the '{concept}' concept. -- Input CV Statement -- {text} -- Instruction -- The candidate's '{concept}' is '{old_value_text}'. Change it to '{new_value_text}'. The following concepts are confounders and must not be changed: {'', '.join(confounders)}'. -- Edited CV Statement -- Return only the updated CV statement with no additional text.</p>
Fix All (Flexible Change)	<p>Box F.15: Fix all Prompt</p> <p>Prompt Instruction: I'm providing a CV statement from the LIBERTy dataset. Your task is to update it by modifying the '{concept}' concept. -- Input CV Statement -- {text} -- Instruction -- The candidate's '{concept}' is '{old_value_text}'. Change it to '{new_value_text}'. The CV statement includes the following concepts: {'', '.join(all_concepts)}'. Some of these concepts are causally linked to '{concept}' and may require adjustments to maintain logical consistency. -- Edited CV Statement -- Return only the updated CV statement with no additional text.</p>
Mediators and Confounders (Causal Framework)	<p>Box F.16: Mediators and Confounders Prompt</p> <p>Prompt Instruction: I'm providing a CV statement from the LIBERTy dataset. Your task is to update it by modifying the '{concept}' concept. -- Input CV Statement -- {text} -- Instruction -- The candidate's '{concept}' is '{old_value_text}'. Change it to '{new_value_text}'. The following concepts are confounders, meaning they must remain unchanged: {'', '.join(confounders)}'. The following concepts are mediators, meaning they are causally influenced by '{concept}' and may require adjustments to maintain logical consistency: {'', '.join(mediators)}'. -- Edited CV Statement -- Return only the updated CV statement with no additional text.</p>

Workplace Violence Prediction												
→ Model ↓ Method	Average		DeBERTa-v3		T5		Qwen-2.5		Llama-3.1		GPT-4o	
	ED	OF	ED	OF	ED	OF	ED	OF	ED	OF	ED	OF
<i>CF Gen</i>	0.47	0.58	0.39	0.71	0.37	0.67	0.54	0.64	0.51	0.32	0.52	0.57
<i>Approx</i>	0.41	0.71	0.34	0.80	0.31	0.76	0.48	0.76	0.42	0.51	0.51	0.70
<i>ConVecs</i>	0.40	0.73	0.27	0.86	0.34	0.77	0.44	0.79	0.42	0.51	0.51	0.71
<i>ST Match</i>	0.51	0.63	0.53	0.68	0.44	0.66	0.65	0.65	0.41	0.50	0.52	0.68
<i>PT Match</i>	0.51	0.64	0.53	0.67	0.44	0.66	0.65	0.65	0.37	0.56	0.57	0.64
<i>FT Match</i>	0.32	0.84	0.11	0.93	0.23	0.83	0.39	0.79	0.38	0.52	0.46	0.72

Disease Detection												
→ Model ↓ Method	Average		DeBERTa-v3		T5		Qwen-2.5		Llama-3.1		GPT-4o	
	ED	OF	ED	OF	ED	OF	ED	OF	ED	OF	ED	OF
<i>CF Gen</i>	0.67	0.36	0.63	0.47	0.54	0.48	0.59	0.46	0.78	0.10	0.79	0.31
<i>Approx</i>	0.48	0.69	0.43	0.74	0.43	0.71	0.51	0.66	0.53	0.63	0.50	0.69
<i>ConVecs</i>	0.44	0.70	0.38	0.74	0.41	0.72	0.46	0.67	0.50	0.63	0.47	0.72
<i>ST Match</i>	0.46	0.69	0.43	0.72	0.41	0.71	0.45	0.68	0.44	0.65	0.56	0.70
<i>PT Match</i>	0.52	0.65	0.49	0.70	0.49	0.66	0.49	0.65	0.49	0.60	0.65	0.66
<i>FT Match</i>	0.36	0.75	0.18	0.86	0.31	0.78	0.39	0.73	0.44	0.66	0.46	0.73
<i>LEACE</i>	0.65	0.46	0.62	0.42	0.46	0.55	0.87	0.41	–	–	–	–

CV Screening												
→ Model ↓ Method	Average		DeBERTa-v3		T5		Qwen-2.5		Llama-3.1		GPT-4o	
	ED	OF	ED	OF	ED	OF	ED	OF	ED	OF	ED	OF
<i>CF Gen</i>	0.52	0.52	0.48	0.58	0.49	0.55	0.73	0.48	0.47	0.39	0.43	0.60
<i>Approx</i>	0.46	0.67	0.36	0.74	0.33	0.71	0.51	0.69	0.50	0.56	0.58	0.63
<i>ConVecs</i>	0.47	0.66	0.38	0.75	0.39	0.71	0.50	0.67	0.52	0.53	0.57	0.62
<i>ST Match</i>	0.50	0.62	0.52	0.67	0.48	0.63	0.56	0.64	0.41	0.56	0.52	0.62
<i>PT Match</i>	0.50	0.63	0.53	0.68	0.49	0.63	0.54	0.65	0.40	0.56	0.55	0.63
<i>FT Match</i>	0.35	0.72	0.19	0.86	0.26	0.78	0.40	0.73	0.42	0.57	0.50	0.65

Table 15: **Local Explainability – Full Results:** We report the Average ICaCE Error-Distance (\overline{ED} , ↓ is better) and Average ICaCE Order-Faithfulness (OF, ↑ is better). The **Average** column reports the mean across five explained models and three datasets.

Workplace Violence Prediction						
→ Model	Average	DeBERTa-v3	T5	Qwen-2.5	Llama-3.1	GPT-4o
<i>CF Gen</i>	0.772	0.810	0.857	0.762	0.476	0.952
<i>Approx</i>	0.781	0.810	0.905	0.810	0.524	0.810
<i>ConVecs</i>	0.829	0.905	0.905	0.952	0.524	0.857
<i>ST Match</i>	0.743	0.714	0.810	0.762	0.571	0.857
<i>PT Match</i>	0.762	0.714	0.810	0.857	0.714	0.714
<i>FT Match</i>	0.857	1.000	1.000	0.762	0.571	0.952
<i>ConceptSHAP</i>	0.444	0.381	0.333	0.619	–	–

Disease Detection						
→ Model	Average	DeBERTa-v3	T5	Qwen-2.5	Llama-3.1	GPT-4o
<i>CF Gen</i>	0.476	0.619	0.333	0.333	0.524	0.571
<i>Approx</i>	0.838	1.000	0.810	0.714	0.810	0.857
<i>ConVecs</i>	0.876	0.905	0.810	0.905	0.857	0.905
<i>ST Match</i>	0.790	0.762	0.714	0.619	0.952	0.905
<i>PT Match</i>	0.629	0.762	0.381	0.524	0.810	0.667
<i>FT Match</i>	0.877	0.905	0.810	0.857	0.905	0.952
<i>LEACE</i>	0.619	0.667	0.571	0.619	–	–
<i>ConceptSHAP</i>	0.333	0.524	0.190	0.286	–	–

CV Screening						
→ Model	Average	DeBERTa-v3	T5	Qwen-2.5	Llama-3.1	GPT-4o
<i>CF Gen</i>	0.599	0.429	0.643	0.464	0.607	0.750
<i>Approx</i>	0.685	0.643	0.750	0.714	0.571	0.750
<i>ConVecs</i>	0.750	0.714	0.786	0.821	0.643	0.786
<i>ST Match</i>	0.650	0.607	0.464	0.643	0.750	0.786
<i>PT Match</i>	0.671	0.607	0.464	0.750	0.750	0.786
<i>FT Match</i>	0.783	0.821	0.857	0.857	0.786	0.643
<i>ConceptSHAP</i>	0.448	0.536	0.500	0.357	–	–

Table 16: **Global Explainability – Full Results:** Mean Order-Faithfulness (\overline{OF} , \uparrow is better) for global explanations of each model and dataset. Bolded values mark the best-performing method per column.

Workplace Violence Prediction								
Model	Race	Gender	Age	Seniority	Department	License	Tenure	—
DeBERT-v3	0.350	1.192	0.758	0.831	0.595	0.595	0.525	—
T5	0.421	0.743	0.512	0.645	0.569	0.452	0.307	—
Qwen-2.5	0.691	1.314	1.045	0.713	1.308	0.656	0.597	—
Llama-3.1	0.224	0.227	0.226	0.208	0.227	0.211	0.211	—
GPT-4o	0.724	0.594	0.300	0.413	1.203	0.279	0.256	—
True Effect	0.484	1.271	1.154	0.560	1.232	0.572	0.613	—

Disease Detection								
Model	Dizzy	Facial Pain	Fever	Weakness	Headache	Nasal Congestion	Light Sensitivity	—
DeBERT-v3	0.505	0.593	0.243	0.415	0.398	0.395	0.693	—
T5	0.352	0.678	0.284	0.376	0.530	0.506	0.745	—
Qwen-2.5	0.495	0.710	0.383	0.512	0.426	0.443	0.679	—
Llama-3.1	0.487	0.442	0.474	0.332	0.364	0.587	0.519	—
GPT-4o	0.364	0.644	0.504	0.215	0.369	0.684	0.879	—

CV Screening								
Model	Race	Gender	Age	Education	Socioeconomic	Volunteering	Experience	Certificates
DeBERT-v3	0.715	0.432	0.613	1.297	0.245	0.391	0.285	0.732
T5	0.742	0.398	0.513	1.143	0.086	0.066	0.168	0.443
Qwen-2.5	0.522	0.361	0.503	0.799	0.354	0.335	0.756	0.425
Llama-3.1	0.374	0.283	0.397	0.437	0.336	0.349	0.381	0.329
GPT-4o	0.417	0.208	0.355	0.679	0.237	0.251	0.727	0.227
True Effect	0.636	0.369	0.913	1.357	0.209	0.586	0.866	0.599

Table 17: **Concept Sensitivity Analysis – Full Results:** We report concept sensitivity as follows: for each example and concept change, we compute ICaCE values and sum their absolute magnitudes across all output classes. The final concept score is then the average of this quantity across all examples and changes. We also report the ground-truth sensitivity of Y from the SCMs, except in the Disease Detection dataset, where Y (the disease) is the parent of the concepts (symptoms).

Guidance

You will now review a personal statement from a job applicant. Your task is to evaluate the statement based on specific criteria to assess the quality of the content.

Please read the personal statement thoroughly. Evaluate the statement based on these criteria:

1. Identify and mark all required personal concepts that are missing from the statement.
2. Rate the fluency, coherence, relevance, and consistency of the text on a scale from 1 to 5.

Assessment

1) Missing Personal Information Identification

Your task is to mark all personal concepts from the given list that are not mentioned in the statement. If a concept is not explicitly stated or cannot be logically inferred, mark it. For 'volunteering' and 'certificates' concepts, follow this rule: if the statement does not mention them at all, assume they do not exist. Therefore: If the concept list says 'Has volunteering', but the statement makes no mention of volunteering, mark it. If the concept list says 'No volunteering', and the statement describes volunteer work, mark it.

- Gender: Female^[1]
- Age: 51^[2]
- Race: Asian^[3]
- Socioeconomic Status: High^[4]
- Education Level: Doctorate degree^[5]
- No volunteering experience^[6]
- Work Experience: 19^[7]
- Has Certificates^[8]

2) Fluency: Assesses the smoothness and grammatical integrity of text, ensuring it reads naturally without errors.

1^[9] 2^[0] 3^[1] 4^[w] 5^[e]

3) Coherence: Measures the logical flow and connectedness of ideas within the text, ensuring clarity and orderly progression.

1^[t] 2^[a] 3^[s] 4^[d] 5^[f]

4) Relevance: Evaluates how closely text relates to the specified topic or context, ensuring it addresses the intended subject matter.

1^[g] 2^[z] 3^[x] 4^[c] 5^[v]

5) Consistency: Checks for alignment of text with itself and external facts, ensuring no contradictions or inaccuracies.

1^[b] 2^[y] 3^[l] 4^[o] 5^[a]

Figure 4: Annotation guidelines for validating concept values and rating coherence, fluency, task relevance, and logical consistency. Example of the CV screening dataset.

Task Guidance

In this task, you will review two texts, labeled Statement A and Statement B. Statement A was written by a person describing their background, skills, and suitability for a role. Statement B shows how that same person might describe themselves in a parallel world, where one or more personal concepts have changed.

Your goal is to assess how likely it is that the same person could have written both statements, assuming only the specified concept(s) were changed.

Use the following 1–5 Likert scale to rate your judgment:

- 1: The person who wrote Statement B could not have written Statement A, even with the concept changes.
- 5: The same person could have written both statements; the only meaningful difference is the specified concept change(s).

Before starting the task, please review the examples below. Each one includes Statement A and Statement B, their corresponding personal concept lists, and a purple highlight indicating which concepts were changed in Statement B. Each example also includes the rating it received to help guide your judgment.

Positive Example

Concepts list for Statement A: {*Education Level: 'Doctorate degree', 'Age: 42, 'Gender: 'Female', 'Race: 'Asian'*}

I'm a 42-year-old Asian woman with a Doctorate in Computer Science, focused on developing innovative algorithmic solutions in the field of machine learning. After returning from maternity leave, I'm now exploring my next professional opportunity.

Concepts list for Statement B: {*Education Level: 'Bachelor degree', 'Age: 42, 'Gender: 'Male', 'Race: 'Asian'*}

The changed concepts are: *Gender Male, Education Level: Bachelor*

As a 42-year-old Asian man with a Bachelor's degree in Computer Science, I focus on building practical software tools to support machine learning efforts. I've recently become open to new opportunities as I consider the next step in my professional journey.

Estimated likelihood that both texts were authored by the same individual in a parallel scenario differing only in the provided concepts list: 5/5

Negative Example

Concepts list for Statement A: {*Education Level: 'Doctorate degree', 'Age: 42, 'Gender: 'Male', 'Race: 'Asian'*}

I am a 42-year-old Asian male with a Doctorate in Computer Science, dedicated to developing ethical and impactful algorithms that support responsible machine learning solutions.

Concepts list for Statement B: {*Education Level: 'Doctorate degree', 'Age: 42, 'Gender: 'Female', 'Race: 'Asian'*}

The changed concepts are: *Gender: Female*

With a Doctorate in Computer Science and over two decades of academic experience, I've mostly worked in support roles on machine learning projects. I'm familiar with core methods but haven't been directly involved in leading initiatives or driving algorithmic innovation.

Estimated likelihood that both texts were authored by the same individual in a parallel scenario differing only in the provided concepts list: 1/5

Estimated likelihood that both texts were authored by the same individual in a parallel scenario differing only in the provided concepts list:

☆☆☆☆☆

Figure 5: Annotation guidelines for rating the plausibility of a text as a genuine counterfactual of the original.