Counterfactual Explanations as Interventions in Latent Space

Riccardo Crupi

Intesa Sanpaolo

Alessandro Castelnovo

Intesa Sanpaolo SpA Milano

daniele regoli (✉ daniele.regoli@sns.it)

Intesa Sanpaolo <https://orcid.org/0000-0003-2711-8343>

Beatriz San Miguel Gonzalez

Fujitsu Laboratories of Europe

Research Article

Keywords: Artificial Intelligence, Machine Learning, Explainable AI, Counterfactual Explanations, Causality, Algorithmic Recourse

Posted Date: August 9th, 2021

DOI: <https://doi.org/10.21203/rs.3.rs-626875/v1>

License: © ⓘ This work is licensed under a Creative Commons Attribution 4.0 International License.
[Read Full License](#)

Counterfactual Explanations as Interventions in Latent Space

Riccardo Crupi · Alessandro Castelnovo · Daniele
Regoli · Beatriz San Miguel Gonzalez

Received: date / accepted: date

Abstract Explainable Artificial Intelligence (XAI) is a set of techniques that allows the understanding of both technical and non-technical aspects of Artificial Intelligence (AI) systems. XAI is crucial to help satisfying the increasingly important demand of *trustworthy* Artificial Intelligence, characterized by fundamental characteristics such as respect of human autonomy, prevention of harm, transparency, accountability, etc. Within XAI techniques, counterfactual explanations aim to provide to end users a set of features (and their corresponding values) that need to be changed in order to achieve a desired outcome. Current approaches rarely take into account the feasibility of actions needed to achieve the proposed explanations, and in particular they fall short of considering the causal impact of such actions. In this paper, we present Counterfactual Explanations as Interventions in Latent Space (CEILS), a methodology to generate counterfactual explanations capturing by design the underlying causal relations from the data, and at the same time to provide feasible recommendations to reach the proposed profile. Moreover, our methodology has the advantage that it can be set on top of existing counterfactuals generator algorithms, thus minimising the complexity of imposing additional causal constraints. We demonstrate the effectiveness of our approach with a set of different experiments using synthetic and real datasets (including a proprietary dataset of the financial domain).

Keywords Artificial Intelligence · Machine Learning · Explainable AI · Counterfactual Explanations · Causality · Algorithmic Recourse

1 Introduction

The use and importance of Artificial Intelligence (AI) systems and, in particular, Machine Learning (ML) models, has increased in many industrial sectors (i.e. finance, healthcare, hiring, transportation,

Riccardo Crupi, Alessandro Castelnovo, Daniele Regoli
Data Science and Artificial Intelligence, Intesa Sanpaolo, Turin, Italy
E-mail: {name.surname}@intesasnpaolo.com

Beatriz San Miguel Gonzalez
Fujitsu Laboratories of Europe, London, UK
E-mail: beatriz.sanmiguelgonzalez@fujitsu.com

etc.), with the purpose, among others, to automate decision-making processes. The majority of these systems have a direct or indirect impact on people's life. Besides the benefits, many ethical concerns have recently emerged with the widespread use of such systems: bias amplification, data privacy, lack of transparency, human oversight, accountability, etc. (Pekka et al., 2018). This is putting increasing pressure on developers and service providers to supply explanations of the models and in particular of their outcomes. The European Commission has recently published a proposal for what is going to be the first attempt ever to insert AI systems and their use in a coherent legal framework (The European Commission, 2021): the proposal devotes a significant attention to the importance of *transparency* of AI systems.

To enhance transparency and trust in AI systems, Explainable Artificial Intelligence (XAI) has become increasingly important (Adadi and Berrada, 2018). In a nutshell, XAI aims to provide information to explain and justify automated results and thus to give the tools to understand the AI system behaviour. The ultimate goals of XAI are many, among which the prevention of possible harms to end users and also the possibility of gaining useful insights to improve the system itself. Therefore, XAI involves the ability to explain both the technical aspects, pertaining to modeling, and the related human decisions. This entails to consider different target audiences of the proposed explanations, such as data scientists, developers, executives, regulatory entities and end users affected by decisions, among others (Arrieta et al., 2020). Against this background, a growing number of methods and approaches have appeared, depending on the type of problem faced and the stakeholder considered (see, e.g., Adadi and Berrada (2018); Arrieta et al. (2020)).

Explainability in Machine Learning (ML) is usually faced either by employing simple and thus intrinsically understandable models, such as linear or logistic regressions, rule-based methods, decision trees, etc., or by using appropriate tools to generate explanations on trained models. These tools are usually categorized as local or global methods, and as model-specific or model-agnostic solutions (Molnar, 2019; Guidotti et al., 2018b). While local methods aim to provide explanations for a single instance and its given outcome, global methods aim to explain the overall behavior of the model. An example of a global method is that of approximating a black-box model with an intrinsically explainable one. Model-specific approaches are tools that can be applied only to some classes of models (e.g. looking at weights of a linear regression), while model-agnostic methods can be used to explain any black-box model. LIME (Ribeiro et al., 2016) and SHAP (Lundberg and Lee, 2017) are two of the most popular model-agnostic local explainers, and they are part of the broader family of methods based on feature contribution, together with other well-known approaches such as Partial Dependence Plots (Friedman, 2001), Accumulated Local Effects (Apley and Zhu, 2020), Individual Conditional Expectation (Goldstein et al., 2015). Other methods are based on local rule extraction, such as Anchors (Ribeiro et al., 2018) and LORE (Guidotti et al., 2018a). These model-agnostic approaches provide explanations by trying to locally approximate the black-box model. On the other hand, example-based methods explain local instances by computing other points in the feature space — the *examples* — with some desired characteristics, such as representing the typical point belonging to some class (prototypes and criticism (Kim et al., 2016; Gurumoorthy et al., 2019)), or being similar to the original point but with enough changes to be given a different outcome (adversarial examples — see Molnar (2019) — and counterfactual explanations (Wachter et al., 2017)).

Counterfactual explanations, first proposed by Wachter et al. (2017), are becoming one of the most popular approaches to explainability in ML within technical, legal and business circles (Barocas et al., 2020). They are local example-based and mostly model-agnostic explanations that construct a set of statements to communicate to end users what could be changed from an original input to receive an alternative decision. Unlike other explainability techniques, this approach imposes no

constraints on model complexity, avoids the disclosure of technical details (protecting trade secrets), provides precise suggestions for achieving a desired outcome and appears to produce explanations that comply with requirements of note-worthy governmental initiatives (Barocas et al., 2020).

The volume of research on counterfactual explanations is growing and different solutions have been proposed (refer to Verma et al. (2020); Stepin et al. (2021) for surveys on counterfactual explanations methods). While these efforts are significant, generally they fall short of generating feasible actions that end users should carry out (Barocas et al., 2020), which is the focus of this work. Obviously, the concept of causality is a key element if we want to find counterfactual explanations that guide end users to act and not only understand the output of a model (Chari et al., 2020). Causal methods can effectively represent cause-effect relationships among variables, thus going in the direction of disentangling the causal effects on the entire system due to direct changes in some of the variables.

In this paper we present CEILS: Counterfactual Explanations as Interventions in Latent Space, a method to generate counterfactual explanations capturing by design the underlying causal relations from the data. This methodology is based on the idea of employing existing counterfactual explanations generators on a latent space that represent the residual degrees of freedom once the causal structure of the problem is taken into account. We demonstrate the effectiveness of our approach with a set of different experiments using synthetic and real datasets (including a proprietary dataset of the financial domain). We evaluate the explanations using a large set of metrics, trying to quantify and pinpoint the key aspects of our proposal. This evaluation is a useful precursor to user studies, where interactions with users and feedback are employed to guide towards the best explanations.

The paper is organised as follows. Section 2 covers the existing background and related work on counterfactual explanations. Section 3 details our proposed methodology, whose main advantages with respect to prior works are highlighted in section 4. Section 5 is devoted to experiment results to evaluate our method, including a detailed definition of the metrics and a discussion of the major findings. Finally, section 6 concludes the paper by summarizing the proposed method and presenting its main limitations and possible directions of improvement.

2 Related work

2.1 Related concepts

Several governmental initiatives towards explainable AI, such as the General Data Protection Regulation (GDPR) in the European Union (The European Union, 2016) and the Defence Advanced Research Projects Agency (DARPA) XAI program of the United States (Gunning and Aha, 2019), as well as the already mentioned European Commission proposal for legislation on AI systems (The European Commission, 2021), endeavour to promote the creation of explanations that can be understood and appropriately trusted by end users. With the goal of approaching user-centric explanations in AI, researchers can use findings from previous research on social science, wherein contrastive and counterfactual explanations are claimed to be inherent to human cognition (Miller, 2019).

In the field of XAI, there seem to be an overlap in the concepts of *contrastive* and *counterfactual* explanations (Stepin et al., 2021). An explanation is *contrastive* when it does not describe the reason for an event to happen (“*Why P?*”), but seek for the reason of an event to occur *relative to another* that did not (“*Why P rather than Q?*”) (Miller, 2019). Counterfactual explanations are defined as a set of statements constructed to communicate what could be changed in the original profile to get a *different* outcome by the decision-making process (Wachter et al., 2017). Therefore, counterfactual

explanations are normally considered contrastive by nature and give a source of valuable complementary information (Byrne, 2019). Indeed, people usually do not ask why a certain prediction was made, but why this prediction was made *instead of another prediction*: therefore, one of the usual requirements for a “good” explanation is precisely to be contrastive (Lipton, 1990; Molnar, 2019). Notice that counterfactual explanations have the additional characteristic of representing a conditional clause (“If X were to occur, then Y would (or might) occur”) (Stepin et al., 2021), thus adding a “causality layer” on the contrastive statement. Indeed both the work of Karimi et al. (2020c) and our proposed methodology present a bridge between “counterfactuals” as intended by causal inference frameworks (Pearl et al., 2016; Spirtes et al., 2000) and “counterfactual explanations” that are usually *not* embedded in formal causal theory frameworks.

Counterfactual explanations are strictly connected to, but different from, *algorithmic recourse*. While the former, as the name suggests, provides an *explanation* of a specific model outcome (by means of showing a scenario as close as possible to the original but reaching a different outcome), the latter provides *recommendations* of what action to undertake in order to gain a different outcome. In layman’s terms, counterfactual explanations inform an individual where they need to get to, but not how to get there (Karimi et al., 2021b). Rephrasing Karimi et al. (2021a), both can be cast in a counterfactual form by asking the following questions:

- *explanation*: what profile would have led to receiving a different outcome?
- *recourse*: what actions would have led me to reach such profile?

Algorithmic recourse refers, in fact, to the set of actions that an individual should perform in order to reach the desired outcome (Joshi et al., 2019; Venkatasubramanian and Alfano, 2020). Notice that the second question somehow incorporates the first one. In other terms, algorithmic recourse is a broad concept, which contains both the counterfactual explanation and the recommendations on how to reach it.

To address the challenge of algorithmic recourse, it is important to distinguish the variables in terms of their level of “actionability” (Karimi et al., 2020c): there are variables that cannot change (e.g. race, sex, date of birth), variables that can change but cannot be directly controlled by the individual (e.g. credit score), and variables that can — at least in principle — be directly acted upon (e.g. bank balance, income, education).

The aforementioned difference between explanations and recourse may seem only a matter of terminology, and indeed in the majority of the literature on counterfactual explanations it is understood that, given a counterfactual observation, it is straightforward to find the set of actions necessary to reach it by simply taking the difference of the two feature vectors (Barocas et al., 2020). But this is true only under very stringent assumptions, that are outlined in Karimi et al. (2020b, 2021a) and that will be made clearer in section 4.

2.2 Generation of explanations

Since the first proposal of counterfactual explanations by Wachter et al. (2017), a large body of research concerning different algorithms and techniques to generate contrastive and counterfactual explanations have been conducted (Verma et al., 2020; Stepin et al., 2021). Most of generation techniques relies on establishing an optimization problem to find the *nearest counterfactual* in the space of features, with respect to the observation to be explained (Wachter et al., 2017; Karimi et al., 2020a; Mohammadi et al., 2020). The metric used to define the distance to be minimized is

sometimes referred to as *proximity*. Moreover, several additional proposals have been put forward to achieve desirable explanatory properties, such as keeping a low number of feature changes (*sparsity*) (Mothilal et al., 2020) or possibly producing more than one counterfactual explanation per each observation, as *diverse* as possible among each other (Mothilal et al., 2020; Karimi et al., 2020a). Van Looveren and Klaise (2019) propose the use of *prototypes* (Kim et al., 2016; Gurumoorthy et al., 2019) to guide the optimization process, with a twofold goal: to find counterfactuals that are “as close as possible” to the distribution of the observed dataset; to speed up the optimization search. Dhurandhar et al. (2018) propose the use of autoencoders trained on the given data to provide explanation that are near the data manifold.

Some works have put forward proposals to embed causality aspects into the counterfactual generation process. For example, Mahajan et al. (2019) introduce a “causal proximity” loss term by which the counterfactual explanation is “pushed” towards regions in which the underlying causality among features is preserved. Karimi et al. (2020c) build on the optimization framework described in Karimi et al. (2020a) and introduce causality by computing counterfactuals through abduction-action-prediction steps as prescribed by Pearl et al. (2016). Karimi et al. (2020b) focus on the case in which only limited causal knowledge is available, thus propose probabilistic approaches (via Gaussian processes or conditional average treatment effect) that allow to relax the assumption of having access to full structural equations.

The counterfactual generation process is usually expressed as an optimization problem, either constrained or unconstrained (see section 3.1 for details) that has been faced with various strategies: gradient-based methods (Wachter et al., 2017; Van Looveren and Klaise, 2019; Mothilal et al., 2020); genetic-based algorithms (Sharma et al., 2019; Guidotti et al., 2018a); graph-based shortest path algorithms (Poyiadzi et al., 2020); by building on formal verification tools and satisfiability modulo theories (SMT) solvers (Karimi et al., 2020a).

Finally, the literature takes into account other aspects as well: Mahajan et al. (2019) report the evaluation of counterfactual explanations with respect to the *computational efficiency* and the amount of time necessary to generate the explanations, which is indeed one of the challenges in this field (Verma et al., 2020). Binns et al. (2018) and Fernandez et al. (2020) evaluate counterfactual explanations in comparison with other XAI approaches, such as feature importance. Miller (2019) reviews relevant papers from disciplines such as philosophy, cognitive science and social psychology, to draw some findings that can be applied to AI.

3 Methodology

3.1 Problem setting

Consider having a ML classifier $\mathcal{C}$ trying to estimate the relationship between a binary target random variable $Y \in \{0, 1\}$ and predictors $X = (X_1, \dots, X_d)$. Typically, one has

$$\hat{Y} = \mathcal{C}(x) = \mathbf{1}_{\{R(x) > t\}}, \quad (1)$$

where $x \in \mathcal{X}$ is a specific realization of X — namely, an observation; $R(x)$ is usually referred to as *score* and is learned to estimate $P(Y = 1 \mid X = x)$; while t is the threshold above which we assign the positive outcome to the observation x .

Given an instance x^0 we want to find a *counterfactual explanation* for x^0 , i.e. a $x^{0,\text{cf}} \in \mathcal{X}$ such that $\mathcal{C}(x^{0,\text{cf}}) \neq \mathcal{C}(x^0)$. Of course, the simple requirement that x^0 and $x^{0,\text{cf}}$ have different outcomes based

on classifier $\mathcal{C}$ (condition that is referred to as *validity* of $x^{0,\text{cf}}$ (Mothilal et al., 2020; Verma et al., 2020)) is a necessary but not sufficient condition to provide “good” counterfactual explanations.

The general formulation can be written as follows (Karimi et al., 2021a):

$$\begin{cases} x^{0,\text{cf}} = \arg \min_{x \in \mathcal{P}_{\mathcal{X}}} \text{dist}(x, x^0), \\ \mathcal{C}(x) \neq \mathcal{C}(x^0); \end{cases} \quad (2)$$

where $\text{dist} : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}^+$ is a suitable distance function over $\mathcal{X}$. The solution of problem (2) provides the *nearest counterfactual explanation* relative to observation x^0 . The space $\mathcal{P}_{\mathcal{X}} \subseteq \mathcal{X}$ is the subset of feature space $\mathcal{X}$ containing *plausible* counterfactuals, i.e. it embodies a set of requirements that $x^{0,\text{cf}}$ should have in order to represent a realistic set of features with respect to the distribution of training data.

Following Karimi et al. (2020c), we shall distinguish between *plausibility* and *feasibility* constraints. Plausibility constraints refer to all the requirements expressed in feature space that go in the direction of having counterfactual explanations realistic with respect to the observed distribution. Feasibility, on the other hand, refers to the fact that a specific counterfactual $x^{0,\text{cf}}$ is actually reachable with a set of actions from the original observation x^0 . The following example will help clarifying the distinction. An individual who is denied a loan may receive a counterfactual explanation where the age is reduced. While this may be perfectly plausible in terms of observed distribution (namely, there are observations in line with the proposed counterfactual) this is definitely not feasible for that individual to reach the proposed counterfactual. This distinction is valuable in particular when discussing *recommendations* besides explanations, i.e. the algorithmic recourse problem (see section 4).

The optimization problem (2) is often relaxed in terms of an unconstrained loss minimization problem (see e.g. Wachter et al. (2017); Van Looveren and Klaise (2019); Karimi et al. (2020a)):

$$x^{0,\text{cf}} = \arg \min_{x \in \mathcal{X}} \left(L_y(x, x^0) + \lambda \text{dist}(x, x^0) + \sum_i \beta_i L_P^i(x, x^0) \right); \quad (3)$$

where we have:

- the L_y term pushing the outcome y corresponding to x away from that of x^0 , i.e. pushing towards the *validity* of $x^{0,\text{cf}}$;
- the dist term keeping x close to x^0 in feature space (*proximity*),
- the L_P^i terms guiding the solution towards *plausible* points in $\mathcal{X}$.

The parameters $\lambda, \{\beta_i\}_i$ control the relative importance of each term. As mentioned in the previous section, several proposals have been put forward for each of the terms in the loss (3), and in particular for the plausibility terms, to have more realistic, and thus useful explanations to be given to end users. For example, Van Looveren and Klaise (2019) propose to use a term that penalizes the distance between x and the nearest prototype of the class other than $\mathcal{C}(x^0)$, while Dhurandhar et al. (2018) add a term penalizing the distance between x and its reconstruction by an autoencoder trained on the given dataset.

We here restrict to the case in which, for a given instance x^0 , a unique counterfactual $x^{0,\text{cf}}$ is found, but it is also reasonable to suggest multiple counterfactuals per observation, which can be done either by simply running the minimization problem (3) multiple times with different random seeds as in Wachter et al. (2017), or changing the formulation (3) to account for direct minimization over

multiple counterfactual candidates, possibly diverse from each other to capture different aspects of the overall explanation, see e.g. Mothilal et al. (2020); Karimi et al. (2020a).

3.2 Counterfactual explanations in latent space

We propose an algorithmic approach that builds on an arbitrary counterfactual explanation optimizer (namely, a strategy for solving a specific formulation of problem (3)) but is able to find counterfactuals taking into account the underlying causal structure by design. In brief, we propose to generate explanations and corresponding recommendations by searching for nearest counterfactuals not in feature space, but in a latent space representing the residual degrees of freedom once the causal structure of the problem at hand is taken into account. This approach has the advantage of providing the end users with feasible actions to reach a desired outcome, and of doing so with — roughly speaking — a “simple” change of variables on top of existing methodologies.

In doing so, we make use of causal graphs and Structural Casual Models (SCM) (which we discuss in more detail in subsections 3.3 and 3.4, respectively). This is in line with what proposed in Karimi et al. (2020c), but expands it in that we provide a methodology to compute causality-based counterfactual explanations without the necessity of a closed form SCM, and that can be used on top of any counterfactual generator, such as the ones proposed in Wachter et al. (2017), Van Looveren and Klaise (2019) or Mothilal et al. (2020). Indeed, other approaches to incorporate causality in counterfactual explanations generators either rely on an *ad hoc* constrained optimization strategy (Karimi et al., 2020c), or they face the problem by adding additional terms in the loss (3) (Mahajan et al., 2019).

The contribution of our proposal is twofold: providing a straightforward way to incorporate causality relationships into the generation of counterfactual explanations for an arbitrary choice of the baseline generator of explanations; providing, besides counterfactual explanations, causal-aware recommendations for algorithmic recourse.

In a nutshell, our proposal can be summed up as:

1. use the SCM to translate the problem from feature space to the space of exogenous and root variables, that we shall call *latent space* hereafter,
2. apply an arbitrary counterfactual explanation optimizer on latent space,
3. translate counterfactuals back to the original feature space.

3.3 Causal graph

Our solution requires to access to a predefined causal graph that encodes the causal relationships among the variables of the dataset. Modeling causal knowledge is complex and challenging since it requires an actual understanding of the relations, beyond statistical evidence. Different causal discovery algorithms have been proposed to identify causal relationships from observational data through automatic methods (Glymour et al., 2019). For example, the Python **Causal Discovery ToolBox** (Kalainathan et al., 2020) includes many existing causal modeling algorithms such as PC (Spirtes et al., 2000), Structural Agnostic Model (SAM) (Kalainathan et al., 2018), Max-Min Parents and Children (MMPC) (Tsamardinos et al., 2003), etc.; the Python library **CausalNex** implements the NOTEARS algorithm by Zheng et al. (2018); the R packages **pcalg** (Kalisch et al.,

2012), `kpcalg` (Verbyla et al., 2017), `bnlearn` (Scutari, 2010) include a vast selection of causal influence algorithms. In general, it is important that domain experts validate the relations detected by the causal discovery routine, or include new ones when deemed necessary. Moreover, experimentation based causal inference is also possible in some specific circumstances, e.g. via randomized experiments, and also by a combination of the observational and experimental methodologies (Mooij et al., 2020).

As usual, we model the underlying causal relationships among features by means of a Directed Acyclic Graph (DAG) $G = (V, E)$, with V set of vertices (or nodes) and E set of directed edges (or links). Nodes of the graph G are composed by the actual variables $X = (X_1, \dots, X_d)$ used as predictors in the model. Moreover, we denote with U exogenous variables, representing factors not accounted for by the features X , and with Y the dependent variable to be predicted/estimated by means of X . In causal graph theory, edges represent not only conditional dependence relations, but are interpreted as the causal impact that the source variable has on the target variable. We refer to nodes with no parents in G as *root nodes*, namely

$$\text{root nodes} = \{v \in G \mid \text{pa}(v) = \emptyset\}.$$

A Structural Causal Model (SCM) is a triplet (X, U, F) where $F : \mathcal{U} \rightarrow \mathcal{X}$ is a set of functions mapping the exogenous (unobserved) variables to the endogenous (observed) ones: $X = F(U)$. These are called Structural Equations (SE) and, besides describing which variables causally impacts which (that is already encoded in the graph G), they also determine *how* these relations work. Therefore, SCM prescriptions are much stronger than simply prescribing a DAG.

The dataset $D_n = \{(x^1, y^1), \dots, (x^n, y^n)\}$ is composed by n i.i.d. realizations of (X, Y) . Each x_i is a d -dimensional vector, each component representing an observed feature. In the same fashion, $\{u^1, \dots, u^n\}$ represent the realizations of the unobserved variables U .

3.4 Structural Equations

Structural Equations are relations describing the precise functional form that links latent variables U to observable ones X . Assuming an Additive Noise Model (ANM) we have the following:

$$X_v = f_v(\text{pa}(X_v)) + U_v, \quad v = 1, \dots, d. \quad (4)$$

Besides this assumption, we also assume *causal sufficiency*, i.e. that there are no counfounders unaccounted for in the specified DAG.

In figure 1 we see an example of DAG representing the causal relationships Karimi et al. (2020c) in the German credit dataset (Bache and Lichman, 2013). Age and gender are root nodes, i.e. they are not caused by any other variable. $U_1, \dots, U_4$ are the latent variables representing all the unobserved external causes. The unobserved causes for root nodes, namely U_1 and U_2 , are indeed redundant, namely $X_{\text{root nodes}} = U_{\text{root nodes}}$. The SE relative to figure 1 have the following form:

$$\text{age} = U_1, \quad (5a)$$

$$\text{gender} = U_2, \quad (5b)$$

$$\text{amount} = f_3(\text{age}, \text{gender}) + U_3, \quad (5c)$$

$$\text{duration} = f_4(\text{amount}) + U_4. \quad (5d)$$

Fig. 1: Causal graph for the German credit dataset

In what follows, instead of specifying/assuming a precise form for each f_v in equation (4), we are going to infer them from observations, namely from the collection $\{x^1, \dots, x^n\}$. Specifically, in the spirit of Pawlowski et al. (2020), we shall learn a regressor model $\mathcal{M}_v$ estimating X_v from $\mathbf{pa}(X_v)$, and then compute the unobserved term as the residual:

$$\hat{U}_v = X_v - \hat{X}_v = X_v - \mathcal{M}_v(\mathbf{pa}(X_v)), \quad v = 1, \dots, d; \quad (6)$$

which is the equivalent of $\hat{U} = \hat{F}^{-1}(X)$, where $\hat{F}$ are the SE estimated by data through $\mathcal{M}_v$ as discussed below.

For root nodes the model $\mathcal{M}_v$ is of course useless, and all the variability is encoded in the latent variable U_v .¹

Since for root nodes r SE simply reduce to $F_r(U) = U_r$, once all models $\mathcal{M}_v$ are learned, following the causal flow in the DAG it is possible to recursively compute the actual function F connecting U to X — namely $X = F(U)$ — by the following relation²:

$$F_j(U) = \mathcal{M}_j \left(\{F_v(U)\}_{v \in \mathbf{pa}(X_j)} \right) + U_j, \quad \forall j \text{ non-root}. \quad (7)$$

The procedure is summed up in Algorithm 1.

3.5 Model in the latent space

Given the model M with score function $R(x)$ as by (1) and Structural Equations $X = F(U)$, it is straightforward to build their composition, effectively obtaining a model estimating Y given U :

$$\mathcal{C}_u(u) = (\mathcal{C} \circ F)(u) = \mathbf{1}_{\{R(F(u)) > t\}}, \quad (8)$$

where $R(F(u))$ is an estimate of $P(Y = 1 \mid U = u)$. Notice that the model $\mathcal{C}_u$ works precisely by following the causal flow of the underlying causal graph and its SCM. Namely, given some realization of the exogenous factors $U = u$, builds the corresponding values for the observed features by recursively applying (7), i.e. by following the causal flow, and then predicts Y by means of the initial model $\mathcal{C}$. Of course, we don't have access to exogenous variables, but this is not an issue

¹ It is just a matter of notation whether to introduce auxiliary latent variables also for root nodes, that we decide to follow for consistency of the formulation.

² With a slight abuse of notation we omit the estimation symbol $\hat{\cdot}$ over F from here on.

when computing counterfactual explanations, since in that case we only need the value for $U = u^0$ corresponding to the instance that we need to explain (x^0) — and possibly the values $\{u^1, \dots, u^n\}$ corresponding to the training dataset $\mathcal{D}$ — and these can be estimated by means of equation (6).

In other words, $\mathcal{C}_u$ is a model that takes in input u in the residual (or latent) space, converts it in the original space $\mathcal{X}$ thanks to the SE $F(u)$ and then predicts its corresponding Y with the model $\mathcal{C}$.

3.6 Counterfactual generation

Once we have obtained the model $\mathcal{C}_u$ relating U and Y , there are three more steps left to obtain the causal counterfactual explanation:

- compute the latent variables $\{u^0, u^1, \dots, u^n\}$,
- generate a counterfactual explanation $u^{0,cf} = u^0 + a$ of the observation u^0 for the model $\mathcal{C}_u$,
- given $u^{0,cf}$, compute the corresponding feature space counterfactual $x^{0,cf} = F(u^{0,cf})$.

The procedure is summed up in Algorithm 2.

Notice that these 3 steps reflect the usual steps of causal counterfactual computation (Pearl et al., 2016): abduction, intervention and prediction.

Abduction is the phase in which the possible events are restricted by the observation of the actual state of the world, namely $X = x^0$. In our framework, computing $U \mid \{X = x^0\}$ is done via equation (6), i.e. as the residuals of regression models $\mathcal{M}_v$. Taking the example of the German dataset (5), we would have:

$$u_1^0 = \text{age}^0, \quad (9a)$$

$$u_2^0 = \text{gender}^0, \quad (9b)$$

$$u_3^0 = f_3(\text{age}^0, \text{gender}^0) - \text{amount}^0, \quad (9c)$$

$$u_4^0 = \text{duration}^0 - f_4(\text{amount}^0), \quad (9d)$$

where it is understood that the f_i 's need to be either known or estimated, e.g. as regressors $\mathcal{M}_i$.

Intervention is the process of acting upon some variables and fixing it to some specific value. In our framework the actual intervention a is computed — via the minimization problem (3) applied to the model $\mathcal{C}_u$ for the observation u^0 — as the minimal shift in latent space needed to reach a different outcome with respect to $\mathcal{C}_u$, namely $u^{0,cf} = u^0 + a$. Being it a shift on exogenous variables, this is actually a form of *soft intervention* (see section 4 for more details).

Prediction step is the moment in which we compute the values of the observed variables X given the latent $U = u^0$ and given the intervention a . In our framework, this is nothing but $x^{0,cf} = F(u^{0,cf})$. Taking again as reference the example of the German dataset (5), (9), we have:

$$\text{age}^{0,cf} = u_1^0 + a_1, \quad (10a)$$

$$\text{gender}^{0,cf} = u_2^0 + a_2, \quad (10b)$$

$$\text{amount}^{0,cf} = f_3(\text{age}^{0,cf}, \text{gender}^{0,cf}) + u_3^0 + a_3, \quad (10c)$$

$$\text{duration}^{0,cf} = f_4(\text{amount}^{0,cf}) + u_4^0 + a_4. \quad (10d)$$

In the next section we discuss in more detail the role of actions and interventions in (causal) counterfactual explanations.

Algorithm 1: Infer Structural Equations $F(U)$

input : Dataset of observed variables $\{x^1, \dots, x^n\}$, causal graph $G = (V, E)$
output: structural equations $F(U)$, $\{\mathcal{M}_v\}_{v \in V}$

// define auxiliary set of nodes for which Structural Equations are already computed, and initialize it with root nodes
 $V_{\text{temp}} = \{v \in V \mid \text{pa}(v) = \emptyset\}$
 $F_v(U) = U_v$ for v root nodes, i.e. such that $\text{pa}(v) = \emptyset$
// loop until SE are computed for all nodes
while $V_{\text{temp}} \neq V$ **do**
 // loop over nodes in G
 for v **in** V **do**
 if $\text{pa}(v) \subseteq V_{\text{temp}}$ and $v \notin V_{\text{temp}}$ **then**
 train a model $\mathcal{M}_v$ estimating X_v from $X_{\text{pa}(v)}$ using data $\{x^1, \dots, x^n\}$
 // build structural equations under additive model hypothesis
 $F_v(U) = \mathcal{M}_v(\{F_j(U)\}_{j \in \text{pa}(v)}) + U_v$
 // add node v into the list of nodes for which SE are already computed
 $V_{\text{temp}} = V_{\text{temp}} \cup \{v\}$
 end
 end
end
Result: structural equations $F_v(U)$, $\mathcal{M}_v$ for all $v \in G$.

Algorithm 2: Train Model $\mathcal{C}$ and generate counterfactuals from residuals

input : dataset of observed variables $\{(x^1, y^1), \dots, (x^n, y^n)\}$, factual observation (x^0, y^0) , structural equations $F(U)$, $\{\mathcal{M}_v\}_{v \in V}$
output: counterfactual satisfying causal constraints $x^{0, \text{cf}}$ and action a

Train a classifier $\mathcal{C}$ with input dataset $\{x^1, \dots, x^n\}$ and target $\{y^1, \dots, y^n\}$.
// Build a model that estimates y 's from exogenous and root nodes
 $\mathcal{C}_u(U) = \mathcal{C}(F(U))$, $\mathcal{C}_u : U \mapsto Y$
// Generate unobserved variables $\{u^0, u^1, \dots, u^n\}$
 $V_{\text{root}} = \{v \in V \mid \text{pa}(v) = \emptyset\}$
 $u_v^i = x_v^i$, $\forall v \in V_{\text{root}}, \forall i = 1, \dots, n$
 $u_v^i = x_v^i - \mathcal{M}_v(\{x_j^i\}_{j \in \text{pa}(v)})$, $\forall v \notin V_{\text{root}}, \forall i = 0, \dots, n$
// Generate explanation in latent space with a counterfactual generator CF
 $u^{0, \text{cf}} = \text{CF}(\mathcal{C}_u, \{(u^i, y^i)\}_{i=1}^n; u^0, y^0)$
// Compute action
 $a = u^{0, \text{cf}} - u^0$
// Compute explanation in feature space
 $x^{0, \text{cf}} = F(u^{0, \text{cf}})$
Result: action a , counterfactual explanation $x^{0, \text{cf}}$

4 Counterfactual explanations and recommended actions

Counterfactual explanations defined as the nearest possible point to the original observation in feature space such that the model outcome is changed, may have problems in terms of realistic *feasibility*. Indeed, while they may result useful in *explaining* the reasons of an outcome for a specific instance — by showing what features should have been different — they may incur in problems when counterfactuals want to be seen as possible *recommendations* on how to take *actions* in order to get a different outcome. Ustun et al. (2019) introduces the framework of *actionable recourse*, which tries precisely to fill the gap between counterfactual explanations and counterfactual recommendations.

Following Karimi et al. (2020c, 2021a), it is useful to draw a line between the notions of *plausibility* and *feasibility* (or *actionability*). As mentioned in section 3.1, we talk about plausibility constraints whenever we refer to conditions *on the feature space* that pertain to have a realistic counterfactual explanation with respect to the training data distribution. Feasibility constraints, instead, are conditions on the *actions* needed to reach some point in feature space. A couple of examples can clarify the apparent redundancy of these two concepts: a person with a low credit rating is unlikely to be granted a loan, thus an intuitive way to provide a counterfactual explanation is to suggest a profile with the same features but higher rating. One problem with this is that it may be an *unrealistic* profile: since there may be other features correlated with rating, and thus the suggested profile may in fact be an outlier for the true distribution. Besides, there is another problem: how can the loan applicant *reach* the suggested profile? Obviously, he cannot force it's rating to be higher: rating can change, but only as a consequence of the changes in other features, and these changes are not prescribed in the suggested profile. Therefore, the suggested profile is neither plausible nor feasible, for two different reasons. Take now the scenario in which a person is denied a loan because he is too old: in this case suggesting to be younger is of course useless since it is not feasible, but the resulting suggested profile would be, in general, perfectly plausible in terms of features.

Depending on the behavior with respect to actions, it is also useful to define (see Karimi et al. (2020c)):

- **immutable** features, as those that cannot change in any way, neither for direct intervention nor for indirect consequences of changes in other variables;
- **mutable but non-actionable** features, as those that can change due to the changes in other connected features, but cannot be directly intervened upon (such as rating in the example above);
- **actionable** features, as the ones that can vary both due to indirect and direct interventions.

Recourse problem is defined by Ustun et al. (2019) as a constrained optimization problem very similar to that of the nearest counterfactual (2):

$$\begin{cases} a^* = \arg \min_{a \in \mathcal{F}_A} \text{cost}(a, x^0), \\ \mathcal{C}(x^0 + a) \neq \mathcal{C}(x^0), \\ x^0 + a \in \mathcal{P}_X; \end{cases} \quad (11)$$

where a^* is the “cheapest” *action* — in terms of the cost function $\text{cost}(a, x)$ — that the individual identified by x^0 needs to perform in order to reach a different model outcome. The space $\mathcal{F}_A \subseteq \mathcal{A}$ is that of all the actions $\mathcal{A}$ that are *feasible*. It is straightforward to notice that, once defined $x = x^0 - a$ and $\text{cost}(a, x^0) = \text{dist}(x, x^0)$, the recourse problem (11) is equivalent to (2) a part for the explicit formulation of feasibility constraints.

The problem with (11) is that it does not take into account the interdependence among variables and the fact that, in general, the change in a variable comes with changes in other variables as well. The natural framework to discuss how this interdependence impacts actions and counterfactual explanations is that of causal reasoning and in particular the notion of *hard* and *soft interventions* (Eberhardt and Scheines, 2007) that we here try to summarize. First of all, notice that the action computed as discussed in section 3.6, namely:

$$a = u^{0,\text{cf}} - u^0,$$

differs for root nodes and non-root nodes. For root nodes there is no difference between the latent variable U_v and the feature X_v , thus in this case the action is simply the difference between the original and the counterfactual values of that feature. For non-root nodes, instead, the situation is different. In general, we can express the relation between the real world features and the action with a straightforward application of SE, namely:

$$x_v^{\text{cf}} - x_v = F(u^{\text{cf}}) - F(u) = f_v(\mathbf{pa}(x_v^{\text{cf}})) - f_v(\mathbf{pa}(x_v)) + a_v, \quad (12)$$

i.e., the change in each feature is the sum of the change as a consequence of its parents' change and an explicit intervention. This is usually called a *soft intervention* (Eberhardt and Scheines, 2007), since the explicit action a_v is performed *in addition* to the changes due to interventions on the ancestors. In contrast, an *hard intervention* is identified by the following formula

$$x_v^{\text{cf}} - x_v = a_v, \quad (13)$$

i.e. when acting on a variable X_v we force it to assume the value $x_v + a_v$, somehow “destroying” or overriding any change due to its ancestors' changes. Therefore, our proposal is actually implementing by default soft interventions. Notice, however, that hard interventions are by far less interesting, and they are in any case much easier to account for: if we want to force an hard intervention on a variable then it is sufficient to cut out all the corresponding incoming edges in the causal graph. Indeed, as mentioned, for root nodes there is no distinction between hard and soft interventions. In other words, hard interventions are simply soft interventions on a modified casual graph.

In general, if we label with $\mathbf{N}, \mathbf{I}$ the set of immutable and actionable features, respectively, we can write the intervention as follows:

$$x_v^{\text{cf}} - x_v = [f_v(\mathbf{pa}(x_v^{\text{cf}})) - f_v(\mathbf{pa}(x_v))] \mathbf{1}_{X_v \notin \mathbf{N}} + a_v \mathbf{1}_{X_v \in \mathbf{I}}. \quad (14)$$

At the practical level, imposing that a variable is mutable but non-actionable is straightforward in our implementation: it simply results in keeping the corresponding latent space variable fixed. On the other hand, imposing that a variable A is immutable is straightforward when it is a root node, while if it has parents we need to consider carefully what *immutable* means in the specific situation: if it means that A must be kept fixed irrespective of the values of its parents, then it is equivalent to perform a null hard intervention, i.e. we need to change the graph by removing its incoming edges and at the same time freeze its latent variable U_A ; if it means that a variable is fixed and that all the variables that can causally impact it should be fixed as well, then we need to treat as non-actionable both A and all its ancestors up to the root level, i.e. we cannot change any of A ancestors, or it would result in possible changes in A . In the experiments we have done (see section 5) we consider immutable features only at the root level, thus the above distinction on the concept of immutability does not apply.

It is now clear that in the standard recourse optimization problem (11), the action definition as $x^{0,\text{cf}} = x^0 + a$ is equivalent to saying that all actionable variables are acted upon via hard interventions, i.e. each action is seen as enforcing a change in the feature overriding any other change due

to variables interdependence. While this may be realistic in some specific cases, it cannot be taken as a paradigm for the general picture. Indeed, a combination of hard and soft interventions is very likely to be the most common situation.

Moreover, even if we assume realistic that all actions are hard interventions, there still may be a causal flow impacting mutable but non-actionable variables. Therefore, neglecting completely the causal structure and sticking to the formulation (11) is equivalent to both assuming that all actions are hard interventions *and* that there are no mutable but non-actionable variables. Indeed, without a causal structure, the distinction between immutable and non-actionable variables loses meaning.

4.1 Computing *ex post* actions of a given counterfactual explanation

We have tried to clarify the reasons why computing algorithmic recourse without taking into account the causal structure of predictors results in a very specific and not much realistic form of interventions. We now try to address the following issue: instead of finding counterfactual explanations via a latent space representation and then computing actions as differences in the latent variables, why don't we find counterfactual explanations x^* via "standard" algorithms, namely via equation (2) or (3), and *then* find the actions that, given the SCM, would lead to that counterfactual profile?

This program is perfectly pursuable by simply computing

$$a = u^* - u^0,$$

where $u^* = F^{-1}(x^*)$ and $u^0 = F^{-1}(x^0)$ are the residuals with respect to the given SCM of the found counterfactual x^* and the original observation x^0 , respectively. We refer to this approach as computing the *ex post* actions, and it would indeed save all the effort to translate the model into the latent space, and it would reduce to the much simpler task of computing residuals. The drawback is that the actions found will not, in general, satisfy feasibility constraints, since causality is here considered only in retrospective, and there's nothing preventing the found counterfactual x^* to be unreachable with respect to underlying causal structure.

A toy example will clarify this: suppose you have two variables A and B , where, e.g., $B = \alpha A + U_B$, with $\alpha > 0$. Namely, an increase in A causes a linear increase in B . Suppose also that, given an observation x^0 , the probability of finding a valid counterfactual is higher in regions where A is higher but B don't. Then, finding x^* will likely get to a situation where A is higher and B is fixed. In terms of actions, this is only possible when B is intervened upon with a *decrease* in order to compensate for the increase caused by A . If, for any reason, there was a feasibility constraint on actions on B such that only increasing interventions are allowed, then the found counterfactual x^* would not correspond to a set of feasible actions.

Therefore, even if it would be in principle possible to compute counterfactual explanations via the standard optimization problem (3) in feature space and *then*, given a SCM, compute the corresponding actions to be used as recommendations, this simpler procedure fails to satisfy, in general, feasibility constraints. Of course there are cases in which this simpler approach results in outcomes very similar to the CEILS outcomes: these are the cases in which feasibility constraints "don't mix with" underlying causality. Indeed, feasibility problems arise when there are constraints on a variable B having mutable parents (A in the example above). In this case, the changes in B due to the changes in its parents — that are not "seen" by the standard counterfactual optimizer — may result in an action on B that is no more compatible with the feasibility constraints. We shall discuss more on this in the next section, devoted to experiments.

Fig. 2: Causal graph relative to the SE (15).

5 Experiments

In this section, we present the experiments conducted on several datasets to validate the CEILS method. We compare our results with a baseline generator of counterfactual explanations using a set of metrics that captures the particularities of our proposal. Next, we describe the datasets used in the experiments and the experiments setup, detailing the definition of the metrics considered. Finally, we discuss the obtained results.

5.1 Datasets

We use for the experiments a synthetic dataset, two public datasets (German Credit and Sachs) and a proprietary dataset of the financial domain.

Synthetic dataset. We generate a toy dataset of 100,000 samples with two features (X_1 and X_2) and a binary outcome (Y) with the following Structural Equations:

$$X_1 = U_1, \quad U_1 \sim \mathcal{N}(-1, 1) \quad (15a)$$

$$X_2 = X_1 + U_2, \quad U_2 \sim \mathcal{N}(5, 1) \quad (15b)$$

$$Y = \mathbf{1}_{(3X_2 - X_1 + U_Y) > t}, \quad U_Y \sim \mathcal{N}(0, 1) \quad (15c)$$

where, in the experiment, t is chosen for simplicity as the median value of $(3X_2 - X_1 + U_Y)$.

Figure 2 depicts the causal graph of the dataset. The rationale behind equations (15) is to have a very simple model that nevertheless allows us to show some crucial aspects of our proposed methodology. In particular, the key feature of equations (15) is to have a non-root node X_2 that has a high impact on the target variable Y . In this way, we expect to have interesting results when considering X_2 a non-actionable feature. To this end, we define two different experiments: first we set X_2 to actionable (Synthetic #1), then we consider X_2 as mutable but non-actionable (Synthetic #2).

German credit dataset (Bache and Lichman, 2013). This dataset contains financial information of 1,000 applicants who are classified into applicants with high and low risk of defaulting on their loans. We consider a subset of the features in the same way as Karimi et al. (2020c). In particular, we use four main features with the causal relations represented in the DAG of figure 1. Moreover, we constrain gender to be immutable and age to increase only.

Fig. 3: Causal graphs employed in the experiments.

Sachs dataset (Sachs et al., 2005). This dataset contains information on protein expression levels in the human immune system. In particular, it consists of 854 observations with 11 independent measurements of phosphorylated molecules derived from immune system cells, subjected to molecular interventions. We base the experiment on the molecules PKC, MEK, Raf and PKA to predict Erk, considering a binary problem with $Y = 1$ when Erk is above the median value, and $Y = 0$ otherwise. The variables are related according to the DAG depicted in figure 3a (obtained from Sachs et al. (2005)). The inhibition and activation of the molecules included in Sachs et al. (2005) define the constraints that we impose in our experiment. In particular, we consider Raf as non-actionable but mutable, we impose that we can act on PKA only by increasing it and on Mek only by decreasing it. As reported in ?, the molecules Raf/Mek/Erk are associated with growth factors and it could be interesting to control Erk, without intervening directly on it, to prevent cell proliferation and apoptosis.

Proprietary dataset. We use a proprietary dataset with 220,304 credit applications (Castelnovo et al., 2020). This contains 8 features, namely gender, age, citizenship, monthly income, bank seniority, requested amount, number of payments and rating, and the information about the granting/non-granting of the loan. The features are related according to the DAG shown in the figure 3b. We refer to Castelnovo et al. (2020) for more details on data and the corresponding causal graph. We consider gender and citizenship as immutable features and age and bank seniority as features that only can increase in value. Moreover, as in the synthetic case, we run two different experiments, one in which rating is set to actionable (Proprietary #1), and one in which it is set to mutable but non-actionable (Proprietary #2).

5.2 Experiments setup

For all the experiments, we model the $\{\mathcal{M}_v\}$ as feed-forward neural networks with 2 hidden layers. The classifier $\mathcal{C}$ is also modeled as a feed-forward neural network with 2 hidden layers with ReLU activation functions. We employ the open source library TensorFlow for the implementations (Abadi et al., 2015).

5.2.1 Baseline generator of counterfactual explanations

We use as baseline generator of counterfactuals the interpretable counterfactual explanations guided by prototypes (Van Looveren and Klaise, 2019) and in particular, the implementation included in the open source library **Alibi** (Klaise et al., 2019). More specifically, we implement a counterfactual generator with loss weights: 0.2 for L_y (**kappa**), 100 for $L_{\text{prototype}}$ (**theta**), 0.5 for the L_1 proximity term (**beta**), and 0.5 for the $L_{k-d \text{ tree}}$ term (**gamma**). We employ the $k-d$ tree term instead of the autoencoder option. We refer to Klaise et al. (2019) for details on these parameters.

We obtain counterfactual explanations first by straight application of the counterfactual generator guided by prototypes (referred to as “Proto” in Table 1), and then by overlying our CEILS procedure.

We evaluate these two sets of counterfactual explanations by means of a collection of metrics discussed in section 5.2.2. These metrics are of course computed only on valid counterfactuals. Notice, however, that the set of valid counterfactual explanations is, in general, dependent on the methodology. Therefore, we also compute the values of the metrics on the *intersection* of valid counterfactual explanations obtained from both methodologies. This is done in order to have a fair comparison, since one of the methods could have, e.g., very good metrics on very few valid counterfactual explanations, namely the ones related to the factual profiles that are the easiest to be explained.

The evaluation is performed on counterfactual explanations generated for 100 random out-of-sample observations for the German credit dataset, 80 for the Sachs dataset, and on 1,000 for the rest of datasets.

5.2.2 Evaluation metrics

Table 1 shows our experimental results with respect to the following set of metrics (see e.g. Verma et al. (2020) and Karimi et al. (2021a) for a review of potential metrics):

- *Validity* is the fraction of generated explanations that are valid counterfactual explanations, i.e. such that $\mathcal{C}(x^{\text{cf}}) \neq \mathcal{C}(x)$. Thus, it reflects the effectiveness of a method in generating explanations.
- *Proximity*, as discussed in section 3.1, measures how far is the counterfactual explanation from the original instance. Following Wachter et al. (2017) and Mothilal et al. (2020), the proximity of continuous features is computed as the mean of feature-wise L_1 distance re-scaled by the Median Absolute Deviation from the median (MAD) of each feature. On the other hand, for categorical features we consider a distance of 1 whenever the counterfactual example differs from the original input:

$$\text{proximity}_{\text{cont}}(x^{\text{cf}}, x) = \frac{1}{n_{\text{cont}}} \sum_{p=1}^{n_{\text{cont}}} \frac{|x_p^{\text{cf}} - x_p|}{\text{MAD}_p}, \quad (16a)$$

$$\text{proximity}_{\text{cat}}(x^{\text{cf}}, x) = \frac{1}{n_{\text{cat}}} \sum_{p=1}^{n_{\text{cat}}} I(x_p^{\text{cf}} \neq x_p); \quad (16b)$$

where n_{cont} and n_{cat} are the number of continuous and categorical features respectively, and MAD_p is the Median Absolute Deviation from the median for the p -th continuous variable.

- *Sparsity* measures the number of features changes that distinguish the counterfactual explanation from the original instance. In particular, to identify relevant perturbations we consider a threshold as in Mothilal et al. (2020): $t = \min(MAD(f), q_{10}(|\tilde{f} - \text{median}(\tilde{f})|))$, $\tilde{f} = [f_i : f_i \neq \text{median}(f)]$

$$\text{sparsity}(x^{\text{cf}}, x) = \sum_{p=1}^d \mathbf{1}_{\{|x_p^{\text{cf}} - x_p| > t_p\}}. \quad (17)$$

Analogously, we compute sparsity also in terms of actions (see below).

- *Distance*, related to the proximity, measures the L_1 distance between counterfactual and factual observations:

$$\text{distance}(x^{\text{cf}}, x) = \|x^{\text{cf}} - x\|_1. \quad (18)$$

All the above metrics refer directly to the counterfactual explanations, thus we refer to them as metrics on feature space. Instead, the following metrics are focused on the evaluation of actions, thus on latent space quantities. Notice that, when considering the baseline method Proto (or, in general, non-causal methods) there is no latent space to be considered, or, equivalently as discussed above, actions are all hard interventions, i.e. shifts in feature space. However, as argued in section 4.1, we could alternatively think of generating counterfactuals with a non-causal method and then compute anyway the corresponding *ex post* action via the SCM. In the section of table 1 devoted to latent space metrics, for rows corresponding to the baseline method we have indeed reported metrics computed with this *ex post* rationale, thus computing the residuals of the generated counterfactual explanations with respect to the SCM.

- *Cost*, as discussed in section 4, is the magnitude of the action needed to reach a counterfactual point. Specifically, we compute the cost as the L_1 norm of the action: $\|a\|_1$. In terms of the feature space, considering the SCM we have:

$$\text{cost}(x^{\text{cf}}, x) = \|F^{-1}(x^{\text{cf}}) - F^{-1}(x)\|_1. \quad (19)$$

Notice that this is in fact equivalent to the *distance* metric, but on the latent space.

- *Feasibility*, as discussed in section 4, pertains to the fact that suggested actions are realistic, i.e. actually doable. Therefore, it includes all the requirements upon actions. For example, when a feature is non-actionable any non-null action on that feature is unfeasible. In the same way, when a feature can only increase (e.g. age), any negative action on that feature would be unfeasible. We measure feasibility as the percentage of explanations whose actions are all feasible. Thus, we establish the following formulation:

$$F^{-1}(x^{\text{cf}})_v = F^{-1}(x)_v \iff a_v = 0, \quad \forall v \text{ non-actionable}; \quad (20a)$$

$$F^{-1}(x^{\text{cf}})_v \geq F^{-1}(x)_v \iff a_v \geq 0, \quad \forall v \text{ increasing only}; \quad (20b)$$

$$F^{-1}(x^{\text{cf}})_v \leq F^{-1}(x)_v \iff a_v \leq 0, \quad \forall v \text{ decreasing only}. \quad (20c)$$

- *Causal plausibility* is inspired by Mahajan et al. (2019), who propose to add this term to the loss in problem (3) in order to keep the generated explanations “as close as possible” to the underlying causal structure. The intuition is to measure, for each feature, the distance between the found

counterfactual observation and the value that the feature should have if perfectly obeying to the SCM, i.e. with null residuals. The idea is to compare x_v^{cf} with $f_v(\mathbf{pa}(x_v^{\text{cf}}))$. To compute this, one has to build the vector

$$w_v = \begin{cases} \mathcal{M}_v(\mathbf{pa}(x_v^{\text{cf}})), & v \text{ non root} \\ x_v^{\text{cf}}, & v \text{ root} \end{cases}$$

then:

$$\text{causal plausibility}(x^{\text{cf}}, x) = \|x^{\text{cf}} - w\|_1. \quad (21)$$

Notice that this is equivalent to computing the L_1 norm of the (non-root) residuals of x^{cf} with respect to the SCM. In other words, this metric measures the distance of the found counterfactual from the profile that satisfies the SCM with zero residuals (except for root nodes).

Moreover, we compute one more value designed not to compare a counterfactual explanation x^{cf} with its factual counterpart x — as the ones introduced above — but rather to directly compare two methods for counterfactual generation. In particular, we are interested in comparing our proposed methodology with the baseline methodology to understand the net impact of our approach on the underlying counterfactual generator engine. To this end, we compute $\|(x^{\text{cf, base}} - x) - a^{\text{CEILS}}\|_1$ — where $(x^{\text{cf, base}} - x)$ is the recommended action by the baseline generator and a^{CEILS} is the action proposed by our methodology — to measure whether the two methods are recommending actions in the same direction. Obviously, this metric can be computed only on valid counterfactual explanations common to both methods.

5.3 Results

Table 1 summarizes the results obtained in the experiments for all the datasets and metrics. As mentioned, we first compute metrics for the explanations obtained with the two methodologies (Proto and Proto + CEILS) and then for the valid explanations common to both methods (grey rows in the Table 1).

Except for validity and feasibility, for each metric we report the median value and the deviation from the median computed over valid counterfactual explanations found by the corresponding method. As mentioned, we include the same computation over valid counterfactuals common to both methods (grey rows in the table).

In what follows, we first describe the results obtained for each dataset, pointing out the main findings, then we summarize them in section 5.4.

Synthetic dataset. For this dataset we run two different experiments, which differ only for the actionability of X_2 . As discussed in section 5.1, X_2 is the feature with highest impact on the target Y , and it is causally dependent on X_1 . Therefore, it is interesting to assess our methodology when X_2 is mutable but non-actionable: CEILS should be able to learn how to employ X_1 in order to impact X_2 and thus Y ; on the other hand, the baseline Proto cannot, this either treats X_2 as actionable, but in this way this suggests the unfeasible action of changing X_2 , or as non-actionable. However, in this case, the baseline Proto method will not be able to move X_2 in any way, hardly resulting in efficient counterfactual generations.

Even if we inspect what actions would have been recommended for the baseline model considering *ex post* actions (see section 4.1), both strategies (X_2 actionable or non-actionable) result in non-feasible suggestions with respect to X_2 : for the actionable case X_2 is changed by the baseline method, but

Table 1: **Results of the experiments.** The table reports the value of nine metrics for six experiments using four datasets (i.e. Synthetic dataset, German credit dataset, Sachs dataset, and Proprietary dataset) to compare the counterfactual explanations generated via prototypes by Van Looveren and Klaise (2019) (Proto) with our proposal on top of the same counterfactual generator (Proto + CEILS). Validity and feasibility are reported as percentages, while for all the other metrics we report the median values and corresponding deviations from the median. Values are computed over standardized feature values, except proximity and sparsity, which involve MAD computation of features. White rows contain values computed on all valid explanations generated by the corresponding method, gray rows display the metrics computed on the intersection of valid explanations of the 2 methods. For the intersection evaluation, in the validity column we report the number of common valid explanations.

	method	validity	feature space				latent space			
			continuous proximity	categorical proximity	sparsity	distance	sparsity on actions	cost	feasibility	causal plausibility
Synthetic #1 (X_2 actionable)	Proto	100%	0.64 $\pm$ 0.49	-	1.34 $\pm$ 0.61	0.65 $\pm$ 0.44	1.34 $\pm$ 0.61	0.65 $\pm$ 0.48	100%	0.43 $\pm$ 0.22
	Proto + CEILS	99%	0.78 $\pm$ 0.61	-	1.48 $\pm$ 0.63	0.88 $\pm$ 0.61	1.47 $\pm$ 0.64	0.73 $\pm$ 0.46	100%	0.33 $\pm$ 0.18
	Proto	#999	0.64 $\pm$ 0.49	-	1.34 $\pm$ 0.61	0.64 $\pm$ 0.44	1.34 $\pm$ 0.61	0.64 $\pm$ 0.48	100%	0.43 $\pm$ 0.22
	Proto + CEILS	#999	0.78 $\pm$ 0.61	-	1.48 $\pm$ 0.63	0.88 $\pm$ 0.61	1.47 $\pm$ 0.64	0.73 $\pm$ 0.46	100%	0.33 $\pm$ 0.18
Synthetic #2 (X_2 non-actionable)	Proto	63.1%	1.4 $\pm$ 0.83	-	1.0 $\pm$ 0.0	1.73 $\pm$ 0.88	2.0 $\pm$ 0.0	2.93 $\pm$ 1.49	0%	0.99 $\pm$ 0.43
	Proto + CEILS	99.5%	1.98 $\pm$ 1.4	-	1.96 $\pm$ 0.23	2.31 $\pm$ 1.29	0.99 $\pm$ 0.08	1.34 $\pm$ 0.76	100%	0.48 $\pm$ 0.28
	Proto	#631	1.4 $\pm$ 0.83	-	1.0 $\pm$ 0.0	1.73 $\pm$ 0.88	2.0 $\pm$ 0.0	2.93 $\pm$ 1.49	0%	0.99 $\pm$ 0.43
	Proto + CEILS	#631	1.14 $\pm$ 0.68	-	1.94 $\pm$ 0.28	1.41 $\pm$ 0.69	0.99 $\pm$ 0.1	0.83 $\pm$ 0.4	100%	0.39 $\pm$ 0.23
German credit	Proto	70%	1.88 $\pm$ 0.57	0.0 $\pm$ 0.0	2.44 $\pm$ 0.67	2.94 $\pm$ 0.53	2.5 $\pm$ 0.63	2.89 $\pm$ 0.52	100%	1.6 $\pm$ 0.13
	Proto + CEILS	77%	1.97 $\pm$ 0.64	0.0 $\pm$ 0.0	2.4 $\pm$ 0.69	2.99 $\pm$ 0.54	2.37 $\pm$ 0.7	2.89 $\pm$ 0.54	100%	1.6 $\pm$ 0.18
	Proto	#67	1.87 $\pm$ 0.58	0.0 $\pm$ 0.0	2.43 $\pm$ 0.68	2.91 $\pm$ 0.51	2.49 $\pm$ 0.64	2.89 $\pm$ 0.52	100%	1.6 $\pm$ 0.12
	Proto + CEILS	#67	1.86 $\pm$ 0.57	0.0 $\pm$ 0.0	2.45 $\pm$ 0.68	2.91 $\pm$ 0.51	2.4 $\pm$ 0.7	2.87 $\pm$ 0.55	100%	1.59 $\pm$ 0.16
Sachs	Proto	37.5%	2.08 $\pm$ 0.75	-	2.1 $\pm$ 0.66	3.07 $\pm$ 0.95	3.63 $\pm$ 0.81	4.13 $\pm$ 0.96	0%	2.36 $\pm$ 0.81
	Proto + CEILS	22.5%	2.35 $\pm$ 0.86	-	3.56 $\pm$ 0.98	3.47 $\pm$ 0.68	1.91 $\pm$ 0.58	2.5 $\pm$ 0.69	100%	2.17 $\pm$ 0.5
	Proto	#18	1.83 $\pm$ 0.73	-	1.83 $\pm$ 0.62	2.64 $\pm$ 0.7	3.44 $\pm$ 0.98	3.93 $\pm$ 0.84	0%	2.77 $\pm$ 0.87
	Proto + CEILS	#18	2.35 $\pm$ 0.86	-	3.56 $\pm$ 0.98	3.47 $\pm$ 0.68	1.89 $\pm$ 0.68	2.5 $\pm$ 0.69	100%	2.17 $\pm$ 0.5
Proprietary #1 (rating actionable)	Proto	100%	1.49 $\pm$ 9.31	0.32 $\pm$ 0.06	1.46 $\pm$ 0.69	1.26 $\pm$ 0.65	1.68 $\pm$ 0.98	1.36 $\pm$ 0.69	98%	3.32 $\pm$ 0.7
	Proto + CEILS	100%	4.12 $\pm$ 17.49	0.32 $\pm$ 0.05	1.65 $\pm$ 1.05	1.39 $\pm$ 0.69	1.43 $\pm$ 0.75	1.26 $\pm$ 0.64	100%	3.33 $\pm$ 0.71
	Proto	#1000	1.49 $\pm$ 9.31	0.32 $\pm$ 0.06	1.46 $\pm$ 0.69	1.26 $\pm$ 0.65	1.68 $\pm$ 0.98	1.36 $\pm$ 0.69	98%	3.32 $\pm$ 0.7
	Proto + CEILS	#1000	4.12 $\pm$ 17.49	0.32 $\pm$ 0.05	1.65 $\pm$ 1.05	1.39 $\pm$ 0.69	1.43 $\pm$ 0.75	1.26 $\pm$ 0.64	100%	3.33 $\pm$ 0.71
Proprietary #2 (rating non-actionable)	Proto	21.8%	289.57 $\pm$ 830.79	0.0 $\pm$ 0.0	2.86 $\pm$ 0.95	2.16 $\pm$ 1.1	3.49 $\pm$ 1.08	2.51 $\pm$ 1.24	6%	3.68 $\pm$ 1.33
	Proto + CEILS	81.6%	474.15 $\pm$ 879.79	0.12 $\pm$ 0.16	3.18 $\pm$ 0.96	8.63 $\pm$ 5.33	2.71 $\pm$ 0.76	8.65 $\pm$ 5.74	100%	10.18 $\pm$ 5.29
	Proto	#218	289.57 $\pm$ 830.79	0.0 $\pm$ 0.0	2.86 $\pm$ 0.95	2.16 $\pm$ 1.1	3.49 $\pm$ 1.08	2.51 $\pm$ 1.24	6%	3.68 $\pm$ 1.33
	Proto + CEILS	#218	43.23 $\pm$ 109.46	0.09 $\pm$ 0.15	2.83 $\pm$ 1.17	1.72 $\pm$ 0.87	2.28 $\pm$ 1.04	1.35 $\pm$ 0.81	100%	3.3 $\pm$ 1.2

very likely in a way that is not compensated by the change in the parent X_1 , resulting in unfeasible net action on X_2 ; for the non-actionable case X_2 is kept fixed by design in the baseline method, but a non-null action in X_2 is unavoidable to keep X_2 fixed while changing X_1 , which is unfeasible.

These considerations are confirmed by the experimental results in table 1 and in the examples shown in table 2.

Indeed the overall results relative to Synthetic #1 and #2 in table 1 show that:

- in terms of feature space metrics the Proto method performs slightly better, and this is in line with the fact that CEILS focuses on *nearest actions* rather than nearest explanations;
- if we compare the CEILS cost with the effort done by the baseline, namely $\|x^{\text{cf}} - x\|_1$ (i.e. the distance metric), then the gain of using CEILS becomes evident in both runs.

Table 2: **Examples of explanations.** Two examples of instances of the Synthetic #2 experiment explained by counterfactual profiles generated via the prototype method of Van Looveren and Klaise (2019) and via CEILS. The Δ stands for the difference between counterfactual and factual instances. The action for the baseline method is the *ex post* action.

	Variable	Factual	counterfactuals		Δ		Action	
			Proto	Proto + CEILS	Proto	Proto + CEILS	Proto	Proto + CEILS
Example 1	Y	0	1	1				
	X ₁	-1.098	-1.364	-0.995	-0.266	0.103	-0.260	0.103
	X ₂	3.896	3.896	4.006	0	0.110	0.242	0
Example 2	Y	1	0	0				
	X ₁	-0.886	-0.726	-0.973	0.159	-0.087	0.162	-0.087
	X ₂	4.080	4.080	3.999	0	-0.081	-0.168	0

- even if we consider the *ex post* action for the baseline Proto method, we can see that for Synthetic #2 (X_2 non-actionable) there is a huge gain in cost, and this is due to the fact that the Proto method pushes X_1 in the “wrong” direction, since it employs causality only *after* computing explanations;
- analogously, feasibility metric confirms our expectations: in the Synthetic #2 experiment the baseline suggest only unfeasible actions.
- causal plausibility, much higher for the baseline in Synthetic #2, confirms once more this evidence.
- in Synthetic #1, instead, both cost, feasibility and causal proximity are comparable: in this case X_2 is actionable, therefore there is no feasibility issue for neither of the two methods.

Table 2 displays two examples of counterfactual explanations for the case in which X_2 is retained non-actionable. If we focus on the first example where $Y = 0 \rightarrow Y = 1$, the Prototype method tends to decrease X_1 (-0.266) because the model $\mathcal{C}$ has learnt the negative dependence of Y in X_1 , and also it can not act on X_2 . However, CEILS “knows” that by decreasing X_1 there is a linear decrease in X_2 , which has a stronger impact on the outcome Y . Indeed, in the action column, we see that CEILS effectively suggests to increase X_1 (0.103) in order to increase the quantity $3X_2 - X_1$, and it does so with less overall effort required from the end user.

Moreover, notice that, in line with expectations, if we compute the SCM *ex post* action for the baseline, then we have non-feasibility for both examples, since there is a non-null action on X_2 . Also in this case, we see that actions on X_1 are done in the “wrong” direction.

Similar arguments can be made for the example 2 included in the table 2 where the counterfactual methods try to modify the target as $Y = 1 \rightarrow Y = 0$.

German credit dataset. As shown in table 1, the results do not present any significant difference among the two methods. This is not surprising, since the only feasibility constraints are on root nodes (gender kept immutable and age not decreasing). Also in terms of effort there is no apparent discrepancy: the distance obtained with the Proto method and the cost of CEILS are almost identical, i.e. there is no gain given by the causal flow. Also in terms of direction, the two methods seem comparable: $\|(x^{\text{cf, base}} - x) - a^{\text{CEILS}}\|_1$ has a median value of 0.16. This may be due to the fact that the causal impact of age on amount and (then) duration is not strong enough to play a significant role. Similarly, *ex post* actions have the same overall cost with respect to CEILS actions.

This experiment shows that the advantage of employing causality on top of standard approaches is not always appreciable, and is highly dependent on the underlying causality structure and on the constraints set over the variables.

Proprietary dataset. For the proprietary dataset, the behavior of the two methods is not dissimilar from the synthetic case, however this experiment involves much more complex causal relationships, and presents a real-world scenario of credit lending. Here, the role of the feature X_2 is played by the feature rating, i.e. a feature extremely important to determine the final outcome, that cannot be controlled directly by the end users, and usually is a complex function of other variables.

Similarly to the synthetic dataset, we run two experiments: Proprietary #1, where we consider rating actionable, and Proprietary #2, where rating is mutable but non-actionable³. The results are in line with the discussion made for the synthetic case, thus we here focus only on some interesting insights peculiar to this case:

- in Proprietary #2 we see that the baseline Proto method is much less efficient than CEILS in providing valid counterfactuals: this is due to the high importance of rating in determining the target variable and to the fact that the Prototype method cannot change rating indirectly, as CEILS does;
- this also explains the odd discrepancy in terms of costs 8.65 for CEILS vs 2.51 for the baseline: this is indeed an artefact of the small number of valid counterfactuals over which this metric is computed for the baseline method, CEILS has higher cost simply because is finding counterfactuals also for “harder” cases. If we compute the metrics on the common valid explanations (grey rows in the table), the situation is reversed.
- In Proprietary #2 we have $\|(x^{\text{cf, base}} - x) - a^{\text{CEILS}}\|_1 = 1.87$, confirming the fact that the two methods suggest recommendations in very different directions.

Table 3 shows an example of explanations generated by both methods in the Proprietary #2 setting. As expected gender and citizenship remain fixed, while age and seniority have equal or higher values with respect to the original instance. The baseline method produces a counterfactual explanation with values far away from the factual profile (i.e. increases the income to 3643.3K and almost doubles the requested amount while decreasing the number of installments). On the other hand, CEILS only suggests to increase the bank seniority and the requested amount⁴. Indeed, increasing the bank seniority results in better rating, which is enough to reach the loan approval. Evidently, an increase in seniority is impossible without a corresponding increase in age: actually, we have treated bank seniority as an actionable feature, but it would have been more appropriate to consider it as mutable only as a consequence of age changes, since seniority cannot be controlled independently of age, or to consider an additional common confounder. Nevertheless, we have decided to keep seniority actionable to focus our discussion on rating and not to limit too much the baseline method (for which it would have been impossible to change seniority as well as rating). Moreover, notice that considering the *ex post* action for the prototype method results in having a net *increase* in rating (i.e worsening). This confirms what discussed in section 4.1, i.e. to keep the rating fixed the non-causal method needs to intervene with a negative action to compensate the change due to the impact of the suggested changes in other features.

Figure 4 shows the distribution of interventions on rating for two scenarios: with rating considered either actionable or non-actionable. First, we can notice that the baseline method has non-null

³ Recall that for methods not accounting for causality, non-actionable is effectively equivalent to immutable, since each feature can change only via direct interventions (see section 4.1).

⁴ Both methods apparently provide the counter-intuitive suggestion of increasing the requested amount: this is due to the fact that the baseline method of Van Looveren and Klaise (2019) searches for explanations as close as possible to the training data distribution. In other words, a too small requested amount would not be plausible with respect to the other suggested features.

Table 3: **Examples of explanations.** Example of an instance of the Proprietary #2 dataset experiment explained by counterfactual profiles generated via the prototype method of Van Looveren and Klaise (2019) and via Proto + CEILS. The Δ stands for the difference between counterfactual and factual instances. The action for the baseline method is the *ex post* action.

Variable	Factual	counterfactuals		Δ		Action	
		Proto	Proto CEILS	Proto	Proto + CEILS	Proto	Proto + CEILS
decision	reject	approve	approve				
gender	1	1	1	0	0	0	0
age	30	30	30	0	0	0	0
citizenship	1	1	1	0	0	0	0
income	56.7K	3700K	56.7K	3643.3K	0	3643.3K	0
seniority	0	8.1	5.4	8.1	5.4	8.1	5.4
amount	210K	409K	320K	199K	110K	199K	110K
installments	48	33.4	53.1	-14.6	5.1	-22.9	0
rating	10	10	9.625	0	-0.375	0.5	0

ex post actions in both cases, thus meaning that non-causal method cannot in any way account for features that are not directly intervened upon but could vary in response to changes in other variables, while CEILS has obviously null actions on rating in the actionable (but mutable) case. Secondly, we see that when rating has no constraints (the actionable case) then the two distributions are very similar: this underlines once again that our approach — and in general considering causality in counterfactual explanations generator — gives results in line with its baseline method when there are no feasibility constraints on important features.

Fig. 4: The distribution of actions on rating for two scenarios: rating considered either actionable or non-actionable (which, for the baseline method, is equivalent to immutable). For the prototype method the *ex post* actions are reported. Values are rescaled to a maximum of 1 for sake of readability.

Sachs dataset. 80 random out-of-sample instances were used to generate counterfactual explanations. These result in 30 valid counterfactuals (37.5%) for the baseline method and 18 (22.5%) for Proto + CEILS. This can be explained by the fact that Proto + CEILS is satisfying the feasibility constraints (i.e. no action on *Raf* and monotonic conditions on *PKA* and *Mek*) and the compliance to the SCM, while Proto alone is ignoring them, when considering *ex post* actions. In other words, the constraints on actions and the compliance to the SCM make, in this case, CEILS less effective in terms of valid feature space explanations.

Table 4: **Examples of explanations.** Example of an instance of the Sachs dataset experiment explained by counterfactual profiles generated via the prototype method of Van Looveren and Klaise (2019) and via Proto + CEILS. The Δ stands for the difference between counterfactual and factual profiles. The action for the baseline method is the *ex post* action.

Variable	Factual	counterfactuals		Δ		Action	
		Proto	Proto + CEILS	Proto	Proto + CEILS	Proto	Proto + CEILS
Y	0	1	1				
Mek	27.9	27.9	6.8	0	-21.1	31.2	0
PKC	10.3	7.0	3.4	-3.3	-6.9	-3.3	-6.9
PKA	319.0	1211.3	850.3	892.3	531.3	881.1	511.2
Raf	47.4	47.4	35.6	0	-11.8	18.0	0

Indeed, the difficulty in providing feasible counterfactuals can be seen in the intersection part. Proximity, sparsity and distance (metrics in feature space) are higher in Proto + CEILS, while sparsity on action, cost and causal plausibility (metrics in latent space) are all worse for the baseline. Note that Proto + CEILS effort (i.e. cost) (2.50) is slightly lower with respect both to Proto distance (2.64) and also to Proto *ex post* cost, meaning that there is a gain in the effort to reach the suggested explanation.

Inspection of table 4 clearly reveals that the change in feature space (Δ) in PKC and PKA is comparable in the two methods, but CEILS, forced to satisfy the non-actionability in Raf and the non-increase in Mek, has an effective *decrease* in Raf and Mek that the baseline has not. In other words, in this experiment we see that feasibility may also be a friction with respect to providing valid counterfactuals, as it could be expected. Moreover, looking at Mek (31.2) and Raf (18.0) values for Proto *ex post* actions, we have once more the confirmation that non-causal methods are not able to account for feasible recommendations with respect to a given SCM.

Notice that the baseline method does have the feasible constraints as well, but they are interpreted as hard interventions — as discussed in 4. Thus, as we see by the example in table 4, Proto keeps Raf and Mek fixed, but it can change the other features independently of these constraints, while CEILS can not, since any change in the variables has impacts on others as by the SCM. In this sense, CEILS is more constrained than its baseline method.

5.4 Discussion

After discussing separately the results of each experiment, we can summarize the overall findings as follows:

- CEILS provides, in general, counterfactual explanations farther in feature space with respect to the baseline method.
- CEILS is almost always more efficient in providing valid counterfactuals. This is more pronounced and noticeable when there are non-actionable constraints. In this case, the baseline method may be not able to provide actual counterfactual explanations or they may be too far to be considered valid, whereas CEILS can indirectly act also on non-actionable features taking into account causal influences.
- There are cases (e.g. Sachs dataset experiment) in which feasibility constraints and SCM compliance may result in a form of friction to find valid counterfactuals, as it could be expected.

- Comparing CEILS and its baseline (non-causal) method in terms of effort to reach the explanation, i.e. confronting their costs (distance in latent space for CEILS and distance in feature space for Proto) almost always results in a better performance for CEILS, again due to its ability in exploiting causal relationships.
- If we take into account the underlying SCM for both approaches, then the baseline method exhibits a very poor performance, and in particular it falls short of suggesting feasible actions to reach valid counterfactuals (as argued in section 4.1).

These findings are completely in line with the fact that CEILS effectively focuses on searching the *nearest counterfactual in latent space*, thus it is optimized to find the less expensive set of actions with respect to an assumed SCM, thus guaranteeing a valid recourse.

6 Conclusions

Against the background of a flourishing literature on Explainable AI and in particular on counterfactual explanations, we have proposed a new approach — Counterfactual Explanations as Interventions in Latent Space (CEILS) — with a twofold goal, namely to take into account causality in generating counterfactual explanations and to employ them to provide feasible recommendations for recourse, while at the same time having the important characteristic of being a methodology easily adaptable on top of existing counterfactual generator engines. The experimental results clearly show that there are cases in which the baseline generator would recommend explanation completely unfeasible with respect to the underlying causal structure, while our approach — on top of the same generator — is able to provide more realistic and reachable counterfactual profiles, often with less effort.

This is a first attempt in the direction of the ambitious target of providing the end user with realistic explanations *and* feasible recommendations to gain the desired output in automatic decision making processes.

As for future work, we will tackle some limitations of our methodology and open challenges in the field of counterfactual explanations. Firstly, it would be important to relax the assumption of having a complete and reliable causal graph, and allow for the possibility of having a causal-aware generator with an underlying *partial* DAG (e.g. Mahajan et al. (2019) discuss this point in their proposal). Secondly, it would be valuable to find methods to relax the assumption of having a completely deterministic SCM in the form of an additive noise model (4) (e.g. Karimi et al. (2020b) take steps in this direction). Another assumption that we should address more properly is that of causal sufficiency, namely the fact the DAGs account for all the common causes of the observed variables, which is indeed a strong requirement and virtually impossible to be validated. Moreover, in our experiments we employ the methodology of Van Looveren and Klaise (2019) as a baseline generator for its remarkable characteristic of guiding the optimization process towards regions of the feature space that are close in distribution with respect to the observed data: we would like to analyze in detail how this entangles with our approach of applying the counterfactual generator in latent space rather than in feature space.

Finally, it would be really useful to embed our proposed methodology in user-interaction tools and perform studies both to validate our method, and also to improve it by taking into account user feedback, possibly allowing the users to change, among other parameters, the feasibility constraints on actions.

Declarations

Acknowledgments. We want to thank Aisha Naseer (Fujitsu Research of Europe), Greta Greco (Intesa Sanpaolo), and Ilaria Penco (Intesa Sanpaolo) for their support in the joint collaboration that has lead to this work.

Funding. Not applicable.

Conflicts of interest. The views and opinions expressed within this paper are those of the authors and do not necessarily reflect the official policy or position of Intesa Sanpaolo or Fujitsu Research of Europe. Assumptions made in the analysis, assessments, methodologies, models and results are not reflective of the position of any entity other than the authors.

Availability of data and material. The data used in the experiments are all publicly available and appropriately referenced, a part from the proprietary dataset, which is instead not available for sharing. Interested researchers may send us direct request to have an anonymized sample of this dataset.

Code availability. All the code for the implementation of the method proposed in this work, as well as for the experiments, will be published in the next future. In the meantime, interested researchers may send us a direct request and we will be happy to share our code.

Authors' contributions. RC conceived and proposed the main research idea. RC and DR worked out the theoretical framework and computations. RC and BSM carried out the implementations and the experiments with the help of DR. BSM and DR wrote the manuscript with support of RC and AC.

References

- Abadi M, Agarwal A, Barham P, Brevdo E, Chen Z, Citro C, Corrado GS, Davis A, Dean J, Devin M, Ghemawat S, Goodfellow I, Harp A, Irving G, Isard M, Jia Y, Jozefowicz R, Kaiser L, Kudlur M, Levenberg J, Mané D, Monga R, Moore S, Murray D, Olah C, Schuster M, Shlens J, Steiner B, Sutskever I, Talwar K, Tucker P, Vanhoucke V, Vasudevan V, Viégas F, Vinyals O, Warden P, Wattenberg M, Wicke M, Yu Y, Zheng X (2015) TensorFlow: Large-scale machine learning on heterogeneous systems. URL <https://www.tensorflow.org/>, software available from tensorflow.org
- Adadi A, Berrada M (2018) Peeking inside the black-box: a survey on explainable artificial intelligence (XAI). IEEE access 6:52138–52160
- Apley DW, Zhu J (2020) Visualizing the effects of predictor variables in black box supervised learning models. Journal of the Royal Statistical Society: Series B (Statistical Methodology) 82(4):1059–1086
- Arrieta AB, Díaz-Rodríguez N, Del Ser J, Bennetot A, Tabik S, Barbado A, García S, Gil-López S, Molina D, Benjamins R, et al. (2020) Explainable Artificial Intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI. Information Fusion 58:82–115
- Bache K, Lichman M (2013) UCI machine learning repository
- Barocas S, Selbst AD, Raghavan M (2020) The hidden assumptions behind counterfactual explanations and principal reasons. In: Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency, pp 80–89

- Binns R, Van Kleek M, Veale M, Lyngs U, Zhao J, Shadbolt N (2018) 'it's reducing a human being to a percentage' perceptions of justice in algorithmic decisions. In: Proceedings of the 2018 Chi conference on human factors in computing systems, pp 1–14
- Byrne RM (2019) Counterfactuals in explainable artificial intelligence (xai): Evidence from human reasoning. In: IJCAI, pp 6276–6282
- Castelnovo A, Crupi R, Del Gamba G, Greco G, Naseer A, Regoli D, Gonzalez BSM (2020) Befair: Addressing fairness in the banking sector. In: 2020 IEEE International Conference on Big Data (Big Data), IEEE, pp 3652–3661
- Chari S, Gruen DM, Seneviratne O, McGuinness DL (2020) Directions for explainable knowledge-enabled systems. arXiv preprint arXiv:200307523
- Dhurandhar A, Chen PY, Luss R, Tu CC, Ting P, Shanmugam K, Das P (2018) Explanations based on the missing: Towards contrastive explanations with pertinent negatives. arXiv preprint arXiv:180207623
- Eberhardt F, Scheines R (2007) Interventions and causal inference. *Philosophy of science* 74(5):981–995
- Fernandez C, Provost F, Han X (2020) Explaining data-driven decisions made by ai systems: the counterfactual approach. arXiv preprint arXiv:200107417 pp 1–33
- Friedman JH (2001) Greedy function approximation: a gradient boosting machine. *Annals of statistics* pp 1189–1232
- Glymour C, Zhang K, Spirtes P (2019) Review of causal discovery methods based on graphical models. *Frontiers in genetics* 10:524
- Goldstein A, Kapelner A, Bleich J, Pitkin E (2015) Peeking inside the black box: Visualizing statistical learning with plots of individual conditional expectation. *Journal of Computational and Graphical Statistics* 24(1):44–65
- Guidotti R, Monreale A, Ruggieri S, Pedreschi D, Turini F, Giannotti F (2018a) Local rule-based explanations of black box decision systems. arXiv preprint arXiv:180510820
- Guidotti R, Monreale A, Ruggieri S, Turini F, Giannotti F, Pedreschi D (2018b) A survey of methods for explaining black box models. *ACM computing surveys (CSUR)* 51(5):1–42
- Gunning D, Aha D (2019) DARPA's explainable artificial intelligence (XAI) program. *AI Magazine* 40(2):44–58
- Gurumoorthy KS, Dhurandhar A, Cecchi G, Aggarwal C (2019) Efficient data representation by selecting prototypes with importance weights. In: 2019 IEEE International Conference on Data Mining (ICDM), IEEE, pp 260–269
- Joshi S, Koyejo O, Vijitbenjaronk W, Kim B, Ghosh J (2019) Towards realistic individual recourse and actionable explanations in black-box decision making systems. arXiv preprint arXiv:190709615
- Kalainathan D, Goudet O, Guyon I, Lopez-Paz D, Sebag M (2018) Sam: Structural agnostic model, causal discovery and penalized adversarial learning
- Kalainathan D, Goudet O, Dutta R (2020) Causal Discovery Toolbox: Uncovering causal relationships in Python. *Journal of Machine Learning Research* 21(37):1–5
- Kalisch M, Mächler M, Colombo D, Maathuis MH, Bühlmann P (2012) Causal inference using graphical models with the R package pcalg. *Journal of Statistical Software* 47:1–26
- Karimi AH, Barthe G, Balle B, Valera I (2020a) Model-agnostic counterfactual explanations for consequential decisions. In: International Conference on Artificial Intelligence and Statistics, PMLR, pp 895–905
- Karimi AH, von Kügelgen J, Schölkopf B, Valera I (2020b) Algorithmic recourse under imperfect causal knowledge: a probabilistic approach. arXiv preprint arXiv:200606831
- Karimi AH, Schölkopf B, Valera I (2020c) Algorithmic recourse: from counterfactual explanations to interventions. arXiv preprint arXiv:200206278

- Karimi AH, Barthe G, Schölkopf B, Valera I (2021a) A survey of algorithmic recourse: contrastive explanations and consequential recommendations. arXiv preprint arXiv:201004050
- Karimi AH, Schölkopf B, Valera I (2021b) Algorithmic recourse: from counterfactual explanations to interventions. In: Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency, pp 353–362
- Kim B, Koyejo O, Khanna R, et al. (2016) Examples are not enough, learn to criticize! criticism for interpretability. In: NIPS, pp 2280–2288
- Klaise J, Van Looveren A, Vacanti G, Coca A (2019) Alibi: Algorithms for monitoring and explaining machine learning models. <https://github.com/SeldonIO/alibi>
- Lipton P (1990) Contrastive explanation. *Royal Institute of Philosophy Supplements* 27:247–266
- Lundberg S, Lee SI (2017) A unified approach to interpreting model predictions. arXiv preprint arXiv:170507874
- Mahajan D, Tan C, Sharma A (2019) Preserving causal constraints in counterfactual explanations for machine learning classifiers. arXiv preprint arXiv:191203277
- Miller T (2019) Explanation in artificial intelligence: Insights from the social sciences. *Artificial intelligence* 267:1–38
- Mohammadi K, Karimi AH, Barthe G, Valera I (2020) Scaling guarantees for nearest counterfactual explanations. arXiv preprint arXiv:201004965
- Molnar C (2019) Interpretable Machine Learning. <https://christophm.github.io/interpretable-ml-book/>
- Mooij JM, Magliacane S, Claassen T (2020) Joint causal inference from multiple contexts. *Journal of Machine Learning Research* pp 1–18
- Mothilal RK, Sharma A, Tan C (2020) Explaining machine learning classifiers through diverse counterfactual explanations. In: Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency, pp 607–617
- Pawlowski N, Castro DC, Glocker B (2020) Deep structural causal models for tractable counterfactual inference. arXiv preprint arXiv:200606485
- Pearl J, Glymour M, Jewell NP (2016) Causal inference in statistics: A primer. John Wiley & Sons
- Pekka A, Bauer W, Bergmann U, Bielíková M, Bonefeld-Dahl C, Bonnet Y, Bouarfa L, et al. (2018) The european commission’s high-level expert group on artificial intelligence: Ethics guidelines for trustworthy ai. Working Document for stakeholders’ consultation Brussels pp 1–37
- Poyiadzi R, Sokol K, Santos-Rodriguez R, De Bie T, Flach P (2020) Face: feasible and actionable counterfactual explanations. In: Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society, pp 344–350
- Ribeiro MT, Singh S, Guestrin C (2016) ” why should i trust you?” explaining the predictions of any classifier. In: Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining, pp 1135–1144
- Ribeiro MT, Singh S, Guestrin C (2018) Anchors: High-precision model-agnostic explanations. In: Proceedings of the AAAI Conference on Artificial Intelligence, vol 32
- Sachs K, Perez O, Pe’er D, Lauffenburger DA, Nolan GP (2005) Causal protein-signaling networks derived from multiparameter single-cell data. *Science* 308(5721):523–529
- Scutari M (2010) Learning bayesian networks with the bnlearn R package. *Journal of Statistical Software* 35(3):1–22, DOI 10.18637/jss.v035.i03
- Sharma S, Henderson J, Ghosh J (2019) Certifai: Counterfactual explanations for robustness, transparency, interpretability, and fairness of artificial intelligence models. arXiv preprint arXiv:190507857
- Spirtes P, Glymour CN, Scheines R, Heckerman D (2000) Causation, prediction, and search. MIT press

- Stepin I, Alonso JM, Catala A, Pereira-Fariña M (2021) A survey of contrastive and counterfactual explanation generation methods for explainable artificial intelligence. *IEEE Access* 9:11974–12001
- The European Commission (2021) Proposal for a Regulation of the European Parliament and of the Council laying down harmonised rules on Artificial Intelligence (Artificial Intelligence Act) and amending certain Union legislative acts. <https://digital-strategy.ec.europa.eu/en/library/proposal-regulation-laying-down-harmonised-rules-artificial-intelligence>
- The European Union (2016) EU General Data Protection Regulation (GDPR): Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation). *Official Journal of the European Union*, <http://data.europa.eu/eli/reg/2016/679/2016-05-04>
- Tsamardinos I, Aliferis CF, Statnikov A (2003) Time and sample efficient discovery of markov blankets and direct causal relations. In: *Proceedings of the ninth ACM SIGKDD international conference on Knowledge discovery and data mining*, pp 673–678
- Ustun B, Spangher A, Liu Y (2019) Actionable recourse in linear classification. In: *Proceedings of the Conference on Fairness, Accountability, and Transparency*, pp 10–19
- Van Looveren A, Klaise J (2019) Interpretable counterfactual explanations guided by prototypes. *arXiv preprint arXiv:190702584*
- Venkatasubramanian S, Alfano M (2020) The philosophical basis of algorithmic recourse. In: *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, pp 284–293
- Verbyla P, Desgranges NIB, Wernisch L (2017) kpcalg: Kernel PC Algorithm for Causal Structure Detection. URL <https://CRAN.R-project.org/package=kpcalg>, r package version 1.0.1
- Verma S, Dickerson J, Hines K (2020) Counterfactual explanations for machine learning: A review. *arXiv preprint arXiv:201010596*
- Wachter S, Mittelstadt B, Russell C (2017) Counterfactual explanations without opening the black box: Automated decisions and the gdpr. *Harv JL & Tech* 31:841
- Zheng X, Aragam B, Ravikumar PK, Xing EP (2018) Dags with no tears: Continuous optimization for structure learning. In: *Advances in Neural Information Processing Systems*, pp 9472–9483