Fair and Adequate Explanations

Nicholas Asher, Soumya Paul, Chris Russell

► To cite this version:

Nicholas Asher, Soumya Paul, Chris Russell. Fair and Adequate Explanations. 5th IFIP TC 5, TC 12, WG 8.4, WG 8.9, WG 12.9 International Cross-Domain Conference for Machine Learning and Knowledge Extraction (CD-MAKE 2021), Aug 2021, Vienna (virtual), Austria. 10.1007/978-3-030-84060-0_6 . hal-03454256

HAL Id: hal-03454256

<https://hal.science/hal-03454256>

Submitted on 29 Nov 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Fair and Adequate Explanations

Nicholas Asher¹, Soumya Paul², and Chris Russell³

¹ IRIT, Université Paul Sabatier, 118 route de Narbonne, 31062 Toulouse, France
asher@irit.fr

² Telindus, 18 rue du Puits Romain, L-8070 Bertrange, Luxembourg
soumya.paul@telindus.lu

³ Amazon Research, Tübingen, Germany
cmruss@amazon.com

Abstract. Recent efforts have uncovered various methods for providing explanations that can help interpret the behavior of machine learning programs. Exact explanations with a rigorous logical foundation provide valid and complete explanations, but they have an epistemological problem: they may be too complex for humans to understand and too expensive to compute even with automated reasoning methods. Interpretability requires *good* explanations that humans can grasp and can compute.

We take an important step toward specifying what good explanations are by analyzing the epistemically accessible and pragmatic aspects of explanations. We characterize sufficiently good, or fair and adequate, explanations in terms of counterfactuals and what we call the *conundra of the explainee*, the agent that requested the explanation. We provide a correspondence between logical and mathematical formulations for counterfactuals to examine the partiality of counterfactual explanations that can hide biases; we define fair and adequate explanations in such a setting. We then provide formal results about the algorithmic complexity of fair and adequate explanations.

1 Introduction

Explaining the predictions of sophisticated machine-learning algorithms is an important issue for the foundations of AI. Recent efforts [34,35,38,19,4] have shown various methods for providing explanations. Among these, model-based, logical approaches that completely characterise one aspect of the decision promise complete and valid explanations.

Such logical methods are thus *a priori* desirable, but they have an epistemological problem: they may be too complex for humans to understand or even to write down in human-readable form. Interpretability requires *epistemically accessible* explanations, explanations humans can grasp *and compute*. Yet what is a sufficiently complete and *adequate* epistemically accessible explanation, a *good explanation* still needs analysis [30]. We propose to characterize sufficiently good, or fair and adequate, explanations in terms of counterfactuals—explanations, that is that are framed in terms of what would have happened had certain conditions (that do not obtain) been the case—and what we call the *conundrum* and

fairness requirements of the *explainee*, the person who requested the explanation or for whom the explanation is intended). It is this conundrum that makes the explainee request an explanation. Counterfactual explanations, as we argue below, are a good place to start for finding accessible explanations, because they are typically more compact than other forms of explanation.

We argue that a fair and adequate explanation is relative to the cognitive constraints and fairness requirements of an explainee $\mathcal{E}$ [5,1,28]. $\mathcal{E}$ asks for an explanation for why π when she wasn't expecting π . Her not expecting π follows from beliefs that must now be revised—how to specify this revision is the conundrum of $\mathcal{E}$. An adequate explanation is a pragmatic act that should solve the conundrum that gave rise to the request for explanation; solving the conundrum makes the explanation useful to $\mathcal{E}$ [15]. In addition, an adequate explanation must lay bare biases that might be unfair or injurious to $\mathcal{E}$ (the fairness constraint). In effect, this pragmatic act is naturally modelled in a game theoretic setting in which the explainer must understand explainee $\mathcal{E}$'s conundrum and respond so as to resolve it. A cooperative explainer will provide an explanation in terms of the type he assigns to $\mathcal{E}$, as the type will encode the relevant portions of $\mathcal{E}$'s cognitive state. On the other hand the explainee will need to interpret the putative explanation in light of her model of the explainer's view of his type. Thus, both explainer and explainee have strategies that exploit information about the other—naturally suggesting a game theoretic framework for analysis.

In developing our view of fair and adequate explanations, we will exploit both the logical theory of counterfactuals [26] and mathematical approaches for adversarial perturbation techniques [24,40,33,23,9,4]. We provide a correspondence between logical and mathematical formulations for counterfactuals, and we analyze how counterfactual explanations can hide biases. We then formalize conundra and fair and adequate explanations, and we develop a game theoretic setting for proving computational complexity results for finding fair and adequate explanations in non cooperative settings.

2 Background on explanations

Following [5,1], we take explanations to be answers to *why* questions. Consider the case where a bank, perhaps using a machine learning program, judges $\mathcal{E}$'s application for a bank loan and $\mathcal{E}$ is turned down. $\mathcal{E}$ is in a position to ask a *why* question like,

- (1) why was I turned down for a loan?

when her beliefs would not have predicted this. Her beliefs might not have been sufficient to infer that she wouldn't get a loan; or her beliefs might have been mistaken—they might have led her to conclude that she would get the loan. In any case, $\mathcal{E}$ must now revise her beliefs to accord with reality. *Counterfactual explanations*, explanations expressed with counterfactual statements, help $\mathcal{E}$ do this by offering an incomplete list of relevant factors that together with unstated properties of $\mathcal{E}$ entail the *explanandum*—the thing $\mathcal{E}$ needs explained, in this case

her not getting the loan. For instance, the bank might return the following answer to (1):

- (2) Your income is €50K per year.
- (3) If your income had been €100K per year, you would have gotten the loan.

The counterfactual statement (3) states what given all of $\mathcal{E}$'s other qualities would have been sufficient to get the loan. But since her income is in fact not €100K per year, the semantics of counterfactuals entails that $\mathcal{E}$ does not get the loan. (3) also proposes to $\mathcal{E}$ how to revise her beliefs to make them accord with reality, in that it suggests that she mistakenly thought that her actual salary was sufficient for getting the loan and that the correct salary level is €100K per year.⁴

Counterfactual explanations, we have seen, are *partial*, because they do not explicitly specify logically sufficient conditions for the prediction. They are also *local*, because their reliance on properties of a particular sample makes them valid typically only for that sample. Had we considered a different individual, say $\mathcal{D}$, the bank's explanation for their treatment of $\mathcal{D}$ might have differed. $\mathcal{D}$ might have had different, relevant properties from $\mathcal{E}$; for instance, $\mathcal{D}$ might be just starting out on a promising career with a salary of €50K per year, while $\mathcal{E}$ is a retiree with a fixed income.

The partiality and locality of counterfactuals make them simpler and more epistemically accessible than other forms of explanation. Moreover, the logical theory of counterfactuals enables us to move from a counterfactual to a complete and logically valid explanation. So in principle counterfactual explanations can provide both rigour and epistemic accessibility. But not just any partiality will do, since partiality makes possible explanations that are misleading, that hide injurious or unfair biases. To show how the partiality of counterfactual explanations can hide unfair biases, consider the following scenario. The counterfactual in (2) might be true but it also might be *misleading*, hiding an unfair bias. (1)-(2) can be true while another, more morally repugnant explanation that hinges on $\mathcal{E}$'s being female is also true. Had $\mathcal{E}$ been male, she would have gotten the loan with her actual salary of €50K per year. A fair and adequate explanation should expose such biases.

We now move to a more abstract setting. Let $\hat{f}: X^n \rightarrow Y$ be a machine learning algorithm, with X^n an n-dimensional feature space encoding data and Y the prediction space. Concretely, we assume that $\hat{f}$ is some sort of classifier. When $\hat{f} = \pi$, an explainee may want an explanation, an answer to the question, "why π ?" We will say that an explanation is an event by an *explainer*, the provider of the explanation, directed towards the explainee (the person requesting the explanation or to whom the explanation is directed) with a conundrum. An explanation will consist of of an *explanandum*, the event or prediction to

⁴ [10] provide a superficially similar picture to the pragmatic one we present, but their aim is rather different, to provide a semantics for argumentation frameworks. For us the pragmatic aspect of explanations is better explained via a game theoretic framework; see below.

be explained, an *explanans*, the information that is linked in some way to the explanandum so as to resolve the explainee’s conundrum. When the explanation is about a particular individual, we call that individual the *focal point* of the explanation.

Explanations have thus several parameters. The first is the scope of the explanation. For a *global* explanation of $\hat{f}$, the explainee wants to know the behavior of $\hat{f}$ over the total space X^n . But such an explanation may be practically uncomputable; and for many purposes, we might only want to know how $\hat{f}$ behaves on a selection of data points of interest or focal points, like $\mathcal{E}$ ’s bank profile in our example.⁵ Explanations that are restricted to focal points are *local* explanations.

Explanations of program behavior also differ as to the nature of the *explanans*. In this paper, we will be concerned with *external* explanations that involve an explanatory link between features of input or feature space X and the output in Y without considering any internal states of the learning mechanism [11]. These are attractive epistemically, because unpacking the algorithms’ internal states and assigning them a meaning can be a very complicated affair.

A third pertinent aspect of explanations concerns the link between *explanans* and the *explanandum*. [14,19,18] postulate a deductive or logical consequence link between *explanans* and *explanandum*. [19] represent $\hat{f}$ as a set of logic formulas $\mathcal{M}(\hat{f})$. By assuming features with binary values⁶, an *instance* is then a set of literals that assigns values to every feature in the feature space. An *abductive* explanation of why π is a subset minimal set of literals $\mathcal{I}$ such that $\mathcal{M}(\hat{f}), \mathcal{I} \models \pi$. Abductive explanations exploits universal generalizations and a deductive consequence relation. They explain why *any instance* $\hat{x}$ that has $\mathcal{I}$ is such that $\hat{f}(\hat{x}) = \pi$ and hence are known as *global* explanations [29].

Counterfactuals offer a natural way to provide epistemically accessible, partial explanations of properties of individuals or focal points. The counterfactual in (3) gives a sufficient reason for $\mathcal{E}$ ’s getting the loan, *all other factors of her situation being equal* or being as equal as possible (*ceteris paribus*) given the assumption of a different salary for $\mathcal{E}$. Such explanations are often called *local* explanations [8,29], as they depend on the nature of the focal point; they are also partial [38], because the antecedent of a counterfactual are not by themselves logically sufficient to yield the formula in the consequent. Deductive explanations, on the other hand, are invariant with respect to the choice of focal point. But because counterfactual explanations exploit *ceteris paribus* conditions, factors that deductive explanations must mention can remain implicit in a counterfactual explanation. Thus, counterfactual explanations are typically more compact and thus in principle easier to understand.⁷ Counterfactuals are also intuitive vehicles for explanations as they also encode an analysis of causation [26].

⁵ We are implicitly assuming that $\hat{f}$ is too complex or opaque for its behaviour to be analyzed statically.

⁶ By increasing the number of literals we can simulate non binary values, so this is not really a limitation as long as the features are finite.

⁷ See [18] for some experimental evidence of this.

2.1 Counterfactual explanations for learning algorithms

The canonical semantics for a counterfactual language $\mathcal{L}$, which is a propositional language to which a two place modal operator $\Box \rightarrow$ is added, as outlined in [26] exploits a possible worlds model for propositional logic, $\mathfrak{A} = \langle W, \leq, \llbracket \cdot \rrbracket \rangle$, where: W is a non-empty set (of worlds), $\leq$ is a ternary similarity relation ($w' \leq_w w''$), and $\llbracket \cdot \rrbracket : P \rightarrow W \rightarrow \{0, 1\}$ assigns to elements in P , the set of proposition letters or atomic formulas of the logic, a function from worlds to truth values or set of possible worlds. Then, where $\models$ represents truth in such a model, we define truth recursively as usual for formulas of ordinary propositional logic and for counterfactuals $\psi \Box \rightarrow \phi$, we have:

Definition 1. $\mathfrak{A}, w \models \psi \Box \rightarrow \phi$ just in case: $\forall w'$, if $\mathfrak{A}, w' \models \psi$ and $\forall w'' (\mathfrak{A}, w'' \models \psi \rightarrow w' \leq_w w'')$, then: $\mathfrak{A}, w' \models \phi$

What motivates this semantics with a similarity relation? We can find both epistemic and metaphysical motivations. Epistemically, finding a closest or most similar world in which the antecedent ϕ of the counterfactual $\phi \Box \rightarrow \psi$ is true to evaluate its consequent ψ follows a principle of belief revision [12], according to which it is rational to make minimal revisions to one’s epistemic state upon acquiring new conflicting information. A metaphysical motivation comes from the link Lewis saw between counterfactuals and causation; $\neg \phi \Box \rightarrow \neg \psi$ implies that if ϕ hadn’t been the case, ψ wouldn’t have been the case, capturing much of the semantics of the statement ϕ *caused* ψ . The truth of such intuitive causal statements, however, relies on the presence of a host of secondary or enabling conditions. Intuitively the statement that if I had dropped this glass on the floor, it would have broken is true; but in order for the consequent to hold after dropping the glass, there are many elements that have to be the same in that counterfactual situation as in the actual world—the floor needs to be hard, there needs to be a gravitational field around the strength of the Earth’s that accelerates the glass towards the floor, and many other conditions. In other words, in order for such ordinary statements to be true, the situation in which one evaluates the consequent of a counterfactual has to resemble very closely the actual world.

Though intuitive, as this logical definition of counterfactuals stands, it is not immediately obvious how to apply it to explanations of learning algorithm behavior. We need to adapt it to a more analytical setting. We will do so by interpreting the similarity relation appealed to in the semantics of counterfactuals as a distance function or norm as in [39] over the feature space X^n , an n -dimensional space, used to describe data points. To fill out our semantics for counterfactuals in this application, we identify instances in X^n as the relevant “worlds” for the semantics of the counterfactuals. We now need to specify a norm for X^n . A very simple norm assumes that each dimension of X^n is orthogonal and has a Boolean set of values; in this case, X^n has a natural L_1 norm or Manhattan or edit distance [36].⁸ While this assumption commits us to the

⁸ In fact, we only assume a finite set of finitely valued features, since an n -valued feature is definable with n Boolean valued features. By complicating the language

fact that the dimensions of X^n capture *all the causally relevant* factors and that they are all independent—both of which are false for typical instances of learning algorithms, it is simple and makes our problem concrete. We will indicate below when our results depend on this simplifying assumption.

A logic of counterfactuals can now exploit the link between logic formulas, features of points in X^n , and a learning algorithm $\hat{f}$ described in [19,22]. Suppose a focal point $\hat{x}$ is such that $\hat{f}(\hat{x}) = \eta$. A counterfactual $A \Box \rightarrow \pi$ that is true at the point $\hat{x}$, where π is a prediction incompatible with η , has an antecedent that is a conjunction of literals, each literal defining a feature value, and that provides a sufficient and minimal shift in the features of $\hat{x}$ to get the prediction π . Each counterfactual that explains the behavior of $\hat{f}$ around a focal point $\hat{x} \in X^n$ thus defines a minimal transformation of the features of $\hat{x}$ to change the prediction. We now define the transformations on X^n that counterfactuals induce.

Definition 2. Let $i \subset n$. A fixed transformation Δ_i is a function $\Delta_i : X^n \rightarrow X^n$ such that for $x \in X^n$, if $\Delta_i(x) = y$, then x and y differ only in the dimensions in i . We write $x =_i x'$ to mean that x and x' share the same values along dimensions i . Given $x \in X^n$, and $\hat{f}(x) = \eta$ and where $\|\cdot\|_{X^n}$ is a natural norm on X^n , we shall be interested in the following types of transformations.

- (i) $\Delta_i(x)$ is **appropriate** if $\hat{f}(\Delta_i(x)) = \pi$ where η and π are two incompatible predictions in Y .
- (ii) $\Delta_i(x)$ is **minimally appropriate** if it is appropriate and in addition, $\forall x' \in X$ such that $\Delta_i(x) =_i x'$ and $\hat{f}(x') = \pi$, $\|x' - x\|_{X^n} \geq \|\Delta_i(x) - x\|_{X^n}$.
- (iii) $\Delta_i(x)$ is **sufficiently appropriate** if it is appropriate and in addition, for any $j \subsetneq i$, $\Delta_j(x)$ is not appropriate.
- (iv) $\Delta_i(x)$ is **sufficiently minimally appropriate** if it is both sufficiently and minimally appropriate.

Note that when X is a space of Boolean features, then conditions (ii) and (iv) of Def. 2 trivially hold. Given a focal point $\hat{x}$ in X^n , minimally appropriate transformations represent the minimal changes necessary to the features of $\hat{x}$ to bring about a change in the value predicted by $\hat{f}$.

Let $\hat{f} : X^n \rightarrow Y$ and consider now a counterfactual language $\mathcal{L}_{\hat{f}}$ with a set of formulas Π that describe the predictions in Y of $\hat{f}$.

Definition 3. A counterfactual model $\mathcal{C}_{X^n, \hat{f}}$ for $\mathcal{L}_{\hat{f}}$ with $\hat{f} : X^n \rightarrow Y$ is a triple $\langle W, \leq, \llbracket \cdot \rrbracket \rangle$ with W a set of worlds $W = X^n$, $\leq$ defined by a norm $\|\cdot\|$ on X^n and $\llbracket \cdot \rrbracket : P \cup \{\Pi\} \rightarrow W \rightarrow \{0, 1\}$ such that for $A \in P$, $\llbracket A \rrbracket_w = 1$ iff w has feature A and for $\pi \in \Pi$, $\llbracket \pi \rrbracket_w = 1$ iff $\hat{f}(w) = \pi$.

Given a counterfactual model $\mathcal{C}_{X^n, \hat{f}}$ for $\mathcal{L}_{\hat{f}}$ with norm $\|\cdot\|$ on X^n , we say that $\|\cdot\|$ is $\mathcal{L}_{\hat{f}}$ **definable** just in case for worlds $w, w_1 \in X^n$, there is a formula ϕ of $\mathcal{L}_{\hat{f}}$ that separates w_1 from all $w_2 \in X^n$ such that $\|w_2 - w\| < \|w_1 - w\|$ —i.e. for all w_2 , $\|w_2 - w\| < \|w_1 - w\|$, $\mathcal{C}_{X^n, \hat{f}}, w_1 \models \phi$ and $\mathcal{C}_{X^n, \hat{f}}, w_2 \not\models \phi$.

and logic [7], we can have probability estimates on literals and so encode continuous feature spaces.

Proposition 1. . Let $\hat{f}: X^n \rightarrow Y$ and let $\mathcal{C}_{X^n, \hat{f}}$ be a counterfactual model for $\mathcal{L}_{\hat{f}}$ with an $\mathcal{L}_{\hat{f}}$ definable norm. Suppose also that $\hat{f}(w) = \eta$. Then:
 $\mathcal{C}_{X^n, \hat{f}}, w \models \phi \Box \rightarrow \pi$, where $\pi \in \Pi$ and ϕ is a separating formula iff there is a minimally appropriate transformation, $\Delta_i: X^n \rightarrow X^n$, where $\hat{f}(\Delta_i(w)) = \pi$, and $\mathcal{C}_{X^n, \hat{f}}, \Delta_i(w) \models A$.

Proposition 1 follows easily from Definitions 1, 2 and 3.

Proposition 1 is general and can apply to many different norms and languages. We will mostly be concerned here with a special and simple case:

Corollary 1 Let $\mathcal{L}_{\hat{f}}$ be a propositional language with a set P of propositional letters, where P is the set of Boolean valued features of X^n , and let $\mathcal{C}_{X^n, \hat{f}}$ be a counterfactual model for $\mathcal{L}_{\hat{f}}$ with an L_1 norm. Then:

$\mathcal{C}_{X^n, \hat{f}}, w \models A \Box \rightarrow \pi$, where $\pi \in \Pi$ and A is a conjunction of literals in P iff there is a minimally appropriate transformation over the dimensions i fixed by A , $\Delta_i: X^n \rightarrow X^n$, where $\hat{f}(\Delta_i(w)) = \pi$, and $\mathcal{C}_{X^n, \hat{f}}, \Delta_i(w) \models A$.

We can generate minimally appropriate transformations via efficient (poly-time) techniques like optimal transport or diffeomorphic deformations [40,33,23,9,4] for computing adversarial perturbations [24]. In effect all of these diverse methods yield counterfactuals or sets of counterfactuals given Proposition 1. A typical definition of an adversarial perturbation of an instance x , given a classifier, is that it is a smallest change to x such that the classification changes. Essentially, this is a counterfactual by a different name. Finding a closest possible world to x such that the classification changes is, under the right choice of distance function, the same as finding the smallest change to x to get the classifier to make a different prediction.⁹

The great advantage of Proposition 1 is that marries efficient techniques to generate counterfactual explanations with the logical semantics of counterfactuals that provides logically valid (LV) explanations from counterfactual explanations, unlike heuristic methods [35,27,35]. Thus, counterfactual explanations build a bridge between logical rigour and computational feasibility.

Proposition 2. A counterfactual explanation given by a minimally appropriate $\Delta_i(\hat{x})$ in $\mathcal{C}_{X^n, \hat{f}}$, with an L_1 norm and X^n with Boolean valued features, yields a minimal, LV explanation in at worst a linear number of calls to an NP oracle.

Proof sketch. The atomic diagram [6] of $\mathcal{C}_{X^n, \hat{f}}$ in which each world is encoded as a conjunction of literals (Boolean values of the features P of X^n together with predictions from Y), encodes $\mathcal{M}(\hat{f})$. Further, given Corollary 1 and Definition 3, each minimally appropriate Δ_i defines a set of literals $\mathcal{L}_{\Delta_i}$ describing $\Delta_i(\hat{x})$ such that $\Delta_i(\hat{x}), \mathcal{M}(\hat{f}) \models \pi$. [19,18] provide an algorithm for finding a subset minimal set of literals $\mathcal{E} \subseteq \mathcal{L}_{\Delta_i}$ with $\mathcal{E}, \mathcal{M}(\hat{f}) \models \pi$ in a linear number relative to $|\mathcal{L}_{\Delta_i}|$ of calls to an NP oracle [21]. $\square$

⁹ Such minimal perturbations may not reflect the ground truth, the causal facts that our machine learning algorithm is supposed to capture with its predictions, as noted by [25]. We deal with this in Section 4.

3 From partial to more complete explanations

We have observed that counterfactual explanations are intuitively simpler than deductive ones, as they typically offer only a partial explanation. In fact there are three sorts of partiality in a counterfactual explanation. First, a counterfactual explanation is *deductively incomplete*; it doesn't specify the *ceteris paribus* conditions and so doesn't specify what is necessary for a proof of the prediction π for a particular focal point. Second, counterfactual explanations are also partial in the sense that they don't specify all the sufficient conditions that lead to π ; they are hence *globally incomplete*. Finally, counterfactuals are partial in a third sense; they are also *locally incomplete*. To explain this sense, we need a notion of *overdetermination*.

Definition 4. A prediction $\pi \in Y$ by $\hat{f} : X \rightarrow Y$ is overdetermined for a focal point $\hat{x} \in X$ if the set of minimally sufficiently appropriate transformations of $\hat{x}$

$$O(\hat{x}, \pi, \hat{f}) = \{\Delta_i : \Delta_i(\hat{x}) \text{ is minimally sufficiently appropriate}\}$$

contains at least two elements.

Locally incomplete explanations via counterfactuals can occur whenever $\hat{f}$'s counterfactual decisions are over-determined for a given focal point. Many real world applications like our bank loan example will have this feature.

Locally incomplete explanations can, given a particular ML model $\mathcal{M}_{\hat{f}}$, hide implicitly defined properties that show $\hat{f}$ to be unacceptably biased in some way and so pose a problem for fair and adequate explanations. Local incompleteness allows for several explanatory counterfactuals with very different *explanans* to be simultaneously true. This means that even with an explanation, $\hat{f}$ may act in ways unknown to the agent $\mathcal{E}$ or the public that is biased or unfair. Worse, the constructor or owner of $\hat{f}$ will be able to conceal this fact if the decision for $\mathcal{E}$ is overdetermined, by offering counterfactual explanations using maps Δ that don't mention the biased feature.

Definition 5. A prejudicial factor P is a map, $P : X^n \rightarrow X^n$ and $\hat{f}$ exhibits a biased dependency on prejudicial factor P just in case for some $i \neq 0$, Δ_i , and for some incompatible predictions η and π ,

$$\hat{f}(\hat{x}) = \hat{f}(\Delta_i(\hat{x})) = \eta \quad \text{and} \quad \hat{f}(P(\hat{x})) = \hat{f}(P(\Delta_i(\hat{x}))) = \pi$$

Dimensions of the feature space that are atomic formulas in $\mathcal{L}_{\hat{f}}$ can provide examples of a prejudicial factor P . But prejudicial factors P may be also implicitly definable in $M_{\hat{f}}$. Assume that $\hat{\cdot}$ is a map from real individuals x to their representation as data points $\hat{x} \in \hat{X}$. Then: P is $M_{\hat{f}}$ implicitly definable just in case: for all x such that $\hat{x} \in \hat{X}$, $x \in \|P\|$ iff for some boolean combination E of atoms of $\mathcal{L}_{\hat{f}}$, $M_{\hat{f}} \models E(\hat{x})$.

We've just described some pitfalls of locally incomplete counterfactual explanations. We now show how to move from a partial picture of the behavior of $\hat{f}$ to a more complete one using counterfactuals. Imagine that at a focal point $\hat{x}$, $\hat{f}(\hat{x}) = \eta$ and we want to know why not π .

Definition 6. In a counterfactual model $\mathcal{C}_{X^n, \hat{f}}$ with a set of Boolean valued features P , the collection of counterfactuals $\mathbf{S}_{\mathcal{C}, \hat{x}, \pi} = \{\phi \Box \rightarrow \pi : \mathcal{C}_{X^n, \hat{f}}, \hat{x} \models \phi \Box \rightarrow \pi \text{ with } \phi \text{ a Boolean combination of values for atoms in } P\}$ true at $\hat{x}$ gives the complete explanation for why π would have occurred at $\hat{x}$.

Appropriate transformations Δ_i on X^n in a counterfactual model $\mathcal{C}_{X^n, \hat{f}}$ to produce π associated with counterfactuals via Proposition 1 can capture $\mathbf{S}_{\mathcal{C}, \hat{x}, \pi}$ and permit us to plot the local complete explanation of $\hat{f}$ around a focal point $\hat{x}$ with regard to prediction π .

Definition 7. $\mathbf{B}_{\mathcal{C}, \hat{x}, \pi} = \{\Delta_i(\hat{x}) : \Delta_i \text{ is a minimal appropriate transformation for some } i \subset n\}$

Proposition 3. In a counterfactual model $\mathcal{C}_{X^n, \hat{f}}$, $\mathbf{B}_{\mathcal{C}, \hat{x}, \pi} = \{y \in X^n : \exists (\phi \Box \rightarrow \psi) \in \mathbf{S}_{\mathcal{C}, \hat{x}, \pi} \text{ such that } y \text{ is a closest } \phi \text{ world to } \hat{x} \text{ where } \mathcal{C}_{X^n, \hat{f}}, y \models \psi\}$.

For the remainder of this section we will fix a counterfactual model $\mathcal{C}_{X^n, \hat{f}}$ to simplify notation.

We are interested in the space $\mathcal{N}_{\hat{f}, \hat{x}, \pi}$ around $\hat{x}$ with boundary $\mathbf{B}_{\hat{x}, \pi}$.

Definition 8. 1. $\mathcal{N}_{\hat{f}, \hat{x}, \pi}$ is the subspace of X^n such that (i) $\hat{x} \in \mathcal{N}_{\hat{f}, \hat{x}, \pi}$ and (ii) $\mathcal{N}_{\hat{f}, \hat{x}, \pi}$ includes in its interior all those points z for which $\hat{f}(z) = \hat{f}(\hat{x})$ and (iii) the boundary of $\mathcal{N}_{\hat{f}, \hat{x}, \pi}$ is given by $\mathbf{B}_{\hat{x}, \pi}$.
2. $\mathcal{N}_{\hat{f}, \pi, \hat{x}}^d$ is a subspace of $\mathcal{N}_{\hat{f}, \hat{x}, \pi}$ with boundary $\mathbf{B}_{\hat{x}, \pi}^d$, where $\mathbf{B}_{\hat{x}, \pi}^d = \mathbf{B}_{\hat{x}, \pi} \cap B_d(\hat{x})$, where $B_d(\hat{x}) = \{y \in X^n : \|y - \hat{x}\| \leq d\}$.
3. $\mathbf{S}_{\hat{x}, \pi}^d = \{y : \exists (\phi \Box \rightarrow \psi) \in \mathbf{S}_{\hat{x}, \pi} \wedge \mathcal{C}_{X^n, \hat{f}}, y \models \psi \wedge \|y - \hat{x}\| \leq d\}$.

The set $\mathbf{S}_{\hat{x}, \pi}$ can have a complex structure in virtue of the presence of *ceteris paribus* assumptions. Because strengthening of the antecedent fails in semantics for counterfactuals, the counterfactuals in (4) relevant to our example of Section 2 are all satisfiable at a world without forcing the antecedents of (4)b or (4)c to be inconsistent:

- (4) a. If I were making €100K euro, I would have gotten the loan.
- b. If I were making €100K or more but were convicted of a serious financial fraud, I would not get the loan.
- c. If I were making €100K or more and were convicted of a serious financial fraud but then the conviction was overturned and I was awarded a medal, I would get the loan.

The closest worlds in which I make €100k do not include a world w in which I make €100k but am also convicted of fraud. Counterfactuals share this property with other conditionals that have been studied in nonmonotonic reasoning [13,32]. However, if the actual world turns out to be like w , then by weak centering (4)a turns out to be false, because the *ceteris paribus* assumption in (4)a is that the actual world is one in which I'm not convicted of fraud.

In $\mathbf{S}_{\hat{x}, \pi}$ we can count how many times the value of the consequent changes as we move from one antecedent to a logically more specific one (e.g., does the

prediction flip from A to $A \wedge C$ or from $A \wedge C$ to $A \wedge C \wedge D$). For generality, we will also include in the number of flips, the flips that happen when we change the Boolean value of a feature—going from A to $\neg A$ for example. We will call the number of flips the *flip degree* of $\mathbf{S}_{\hat{x},\pi}$.

There is an important connection between the flip degree of $\mathbf{S}_{\hat{x},\pi}$ and the geometry of $\mathcal{N}_{\hat{f},\hat{x},\pi}$. In a counterfactual model, the move from one antecedent ϕ_1 of a counterfactual c_1 to a logically more specific antecedent ϕ_2 of c_2 , with $c_1, c_2 \in \mathbf{S}_{\hat{x},\pi}$ will, given certain assumptions about the underlying norm yield $\hat{x} < y < z$, with y being a closest to $\hat{x}$ point verifying ϕ_1 and z a closest point verifying ϕ_2 . In fact we generalize this property of norms.

Definition 9. A norm $\|\cdot\|$ in a counterfactual model $\mathcal{C}_{X^n,\hat{f}}$ respects the logical specificity of the model iff for any $z \in X^n$ such that $\mathcal{C}_{X^n,\hat{f}}, z \models \psi$ and for counterfactual antecedents $\phi_1, \phi_2, \dots, \phi_n$ describing features of X^n such that $\mathcal{C}_{X^n,\hat{f}}, z \models \phi_1 \square \rightarrow \neg\psi, \phi_2 \square \rightarrow \psi, \dots, \phi_n \square \rightarrow \neg\psi$ such that $\phi_{i+1} \models \phi_i$ and $\phi_i \not\models \phi_{i+1}$, there are collinear $x_1, \dots, x_n \in X^n$ such that for each i , x_i is a closest point to z such that $\mathcal{C}_{X^n,\hat{f}}, x_i \models \phi_i$ and $\|x_{i+1} - z\| > \|x_i - z\|$.

Remark 1 An $L1$ norm for a counterfactual model is a logical specificity respecting norm.

In addition, a flip (move from a point verifying ϕ_1 to a point verifying ϕ_2) corresponds to a move from a transformation Δ_i to a transformation Δ_j with $i \subset j$. Thus, flips determine a partial ordering under $\subseteq$ over the shifted dimensions i : thus $\Delta_i \leq \Delta_j$, if $i \subseteq j$. We are interested in the behavior of $\hat{f}$ with respect to the partial ordering on Δ_i .

Definition 10. $\hat{f}$ is nearly constant around $\hat{x}$, if for every sufficiently minimally appropriate Δ_i for all $\Delta_j \supset \Delta_i$, $\hat{f}(\Delta_j(\hat{x})) = \hat{f}(\Delta_i(\hat{x}))$.

A nearly constant $\hat{f}$ changes values only once for each combination of features/dimensions d_i moving out from a focal point $\hat{x}$. So at some distance d , nearly constant $\hat{f}$ becomes constant $\hat{f}$. For a nearly constant $\hat{f}$ around $\hat{x}$, $\mathbf{S}_{\hat{x},\pi}$, has flip degree 1. A complete local explanation for $\hat{f}$'s prediction of π within d , $\mathbf{S}_{\hat{x},\pi}^d$, is a global explanation $\hat{f}$'s behavior with respect to π .

We can generalize this notion to define an n -shifting $\hat{f}$. If $\hat{f}$ flips values n times moving out from $\hat{x}$, $\mathbf{S}_{\hat{x},\pi}$ has flip degree n .

Proposition 4. Suppose A counterfactual model has a logical specificity respecting norm, then: $\mathbf{S}_{\hat{x},\pi}$, has a flip degree ≤ 2 iff $\mathcal{N}_{\hat{f},\hat{x},\pi}$ forms a convex subspace of $\hat{f}[X]$.

Proof sketch. Assume $\mathbf{S}_{\hat{x},\pi}$ has flip degree ≥ 3 . Then $\mathbf{S}_{\hat{x},\pi}$ will contain counterfactuals with antecedents ϕ, χ, δ such that $\phi \models \chi \models \delta$ but, say, ϕ and δ counterfactually support π but not χ . As the underlying norm respects $\models$, there are collinear points x, y , and z , where x is a closest point to $\hat{x}$ where ϕ is true,

y is a closest χ world, and z is a closest δ world such that $\hat{x} < z < y < x$. But $\hat{x}, y_\chi \in \mathcal{N}_{\hat{f}, \hat{x}, \pi}$, while $x_\phi, z_\delta \in \mathbf{B}_{\hat{x}, \pi}$ and $\notin \mathcal{N}_{\hat{f}, \hat{x}, \pi}$, which makes $\mathcal{N}_{\hat{f}, \hat{x}, \pi}$ non convex. Conversely, suppose $\mathcal{N}_{\hat{f}, \hat{x}, \pi}$ is non convex. Using the construction of counterfactuals from the boundary $\mathbf{B}_{\hat{x}, \pi}$ of $\mathcal{N}_{\hat{f}, \hat{x}, \pi}$ will yield a set with flip degree 3 or higher. $\square$

The flip degree of $\mathbf{S}_{\hat{x}, \pi}$ gives a measure of the degree of non-convexity of $\mathcal{N}_{\hat{f}, \hat{x}, \pi}$, and a measure of the complexity of an explanation of $\hat{f}$'s behavior. A low flip degree for $\mathbf{S}_{\hat{x}, \pi}^d$ with minimal overdeterminations provides a more general and comprehensive explanation. With Proposition 4, a low flip degree converts a local complete explanation into a global explanation, which is *a priori* preferable. It is also arguably closer to our prior beliefs about basic causal processes. The size of $\mathbf{S}_{\hat{x}, \pi}^d$ gives us a measure to evaluate $\hat{f}$ itself; a large $\mathbf{S}_{\hat{x}, \pi}^d$ doesn't approximate very well a good scientific theory or the causal structures postulated by science. Such a $\hat{f}$ lacks generality; it has neither captured the sufficient nor the necessary conditions for its predictions in a clear way. This could be due to a bad choice of features determining $\hat{f}$'s input X^n [9]; too low level or unintuitive features could lead to lack of generality with high flip degrees and numerous overdeterminations. Thus, we can use $\mathbf{S}_{\hat{x}, \pi}^d$ to evaluate $\hat{f}$ and its input representation X^n .

The flip degree of $\mathbf{S}_{\hat{x}, \pi}$ and the topology of $\mathcal{N}_{\hat{f}, \hat{x}, \pi}$ can also tell us about the relation between counterfactual explanations based on some element in X and ground truth instances provided during training. Our learning algorithm $\hat{f}$ is trying to approximate or learn some phenomenon, which we can represent as a function $f : X \rightarrow Y$; the observed pairs $(z, f(z))$ are ground truth points for $\hat{f}$. Ideally, $\hat{f}$ should fit and converge to f —i.e., with the number of data points N $\hat{f}$ is trained on $\lim_{N \rightarrow \infty} \hat{f}^N \rightarrow f$; in the limit explanations of the behavior of $\hat{f}$ will explain f , the phenomenon we want to understand. Given that we generate counterfactual situations using techniques used to find adversarial examples, however, counterfactual explanations may also be based on adversarial examples that have little to no intuitive connection with the ground truth instances $\hat{f}$ was trained on. While these can serve to explain the behavior of $\hat{f}$ and as such can be valuable, they typically aren't good explanations of the phenomenon f that $\hat{f}$ is trying to model. [25] seek to isolate good explanations of f from the behavior of $\hat{f}$ and propose a criterion of topological connectedness for good counterfactual explanations. This idea readily be implemented as a constraint on $\mathcal{N}_{\hat{f}, \hat{x}, \pi}$: roughly, $\hat{f}$ as an approximation of f will yield good counterfactual explanations relative to a focal point $\hat{x}$ only if for any point y outside of $\mathcal{N}_{\hat{f}, \hat{x}, \pi}$, there is a region C where $\hat{f}$ returns the same value and a path of points $y_1, \dots, y_n \in C$ between y and a ground truth data point p such that $f(p) = \hat{f}(p) = \hat{f}(y_i) = \hat{f}(y)$.¹⁰

¹⁰ We note that our discussion and constraint make clear the distinction between f and $\hat{f}$ which is implicit in [25, 15].

4 Pragmatic constraints on explanations

While we have clarified the partiality of counterfactual explanations, AI applications can encode data via hundreds even thousands of features. Even for our simple running example of a bank loan program, the number of parameters might provide a substantial set of counterfactuals in the complete local explanation given by $\mathbf{S}_{\hat{x}, \pi}$. This complete local explanation might very well involve too many counterfactuals for humans to grasp. We still to understand what counterfactual explanations are *pragmatically relevant* in a given case.

Pragmatic relevance relies on two observations. First, once we move out a certain distance from the focal point, then the counterfactual shifts intuitively cease to be about the focal point; they cease to be counterparts of $\hat{x}$ and become a different case. Exactly what that distance is, however, will depend on a variety of factors about the explainee $\mathcal{E}$ and what the explainer believes about $\mathcal{E}$. Second, appropriate explanations must respond to the particular *conundrum* or cognitive problem that led $\mathcal{E}$ to ask for the explanation [5,1,28]. On our view, the explainee $\mathcal{E}$ requires an explanation when her beliefs do not lead her to expect the observed prediction π . When $\mathcal{E}$'s beliefs suffice to predict $\hat{f}(\hat{x}) = \pi$, she has *a priori* an answer to the question *Why did $\hat{f}(x) = \pi$?* In our bank example from Section 2, had $\mathcal{E}$'s beliefs been such that she did not expect a loan from the bank, she wouldn't have needed to ask, *why did the bank not give me a loan?*¹¹

The conundrum comes from a mismatch between $\mathcal{E}$'s understanding of what $\hat{f}$ was supposed to model (our function f) and $\hat{f}$'s actual predictions. So $\mathcal{E}$, in requesting an explanation of $\hat{f}$'s behavior, might also want an explanation of f itself (see the previous section for a discussion). Either $\mathcal{E}$ is mistaken about the nature of $\hat{f}$, or her grasp of $\hat{f}$ is incomplete.¹² More often than not, $\mathcal{E}$ will have certain preconceptions about $\hat{f}$, and then many if not most of the counterfactuals in $\mathbf{S}_{\hat{x}, \pi}$ may be irrelevant to $\mathcal{E}$. A relevant or fair and adequate explanation for $\mathcal{E}$ should provide a set $\mathfrak{C}_{\mathcal{E}}^d$ of appropriate Δ_i with $\|\Delta_i(\hat{x}) - \hat{x}\| \leq d$ showing which of $\mathcal{E}$'s assumptions were faulty or incomplete, thus solving her conundrum.

Suppose that the explainee $\mathcal{E}$ requests an explanation why $\hat{f}(\hat{x}) = \eta$, and that $\hat{x}$ is decomposed into $\langle x_{d_1}, x_{d_2} \rangle$.

CI Suppose $\mathcal{E}$'s conundrum based on incompleteness; i.e., the conundrum arises from the fact that for $\mathcal{E}$ $\hat{f}$ only pays attention to the values of dimensions d_1 in the sense that for her $\hat{f}(\langle x_{d_1}, x_{d_2} \rangle) = \hat{f}(\langle x_{d_1}, x'_{d_2} \rangle)$, for any values x'_{d_2} . Then there is a $\Delta \in \mathfrak{C}_{\mathcal{E}}^d$ such that $\Delta(\langle x_{d_1}, x_{d_2} \rangle) = \langle x_{d_1}, y_{d_2} \rangle$ and $\hat{f}(\Delta(\hat{x})) = \hat{f}(\langle x_{d_1}, y_{d_2} \rangle) = \pi$ while $\hat{f}(\hat{x}) = \hat{f}(\langle x_{d_1}, x_{d_2} \rangle) = \eta$.

CM Suppose $\mathcal{E}$'s conundrum is based on a mistake. Then there is a $\Delta \in \mathfrak{C}_{\mathcal{E}}^d$ such that $\Delta(\langle x_{d_1}, x_{d_2} \rangle) = \langle y_{d_1}, x_{d_2} \rangle$ such that $\hat{f}(\langle y_{d_1}, x_{d_2} \rangle) = \hat{f}(\Delta(\hat{x})) = \pi$. I.e.,

¹¹ Of course $\mathcal{E}$ might want to know whether her beliefs matched the bank's reasons for denying her a loan, but that's a different question—and in particular it's not a *why* question.

¹² Perhaps $\mathcal{E}$ is also mistaken about or has an incomplete grasp of f or if not, she is mistaken about how $\hat{f}$ differs from f . But we will not pursue this here.

Δ must resolve $\mathcal{E}$'s conundrum by providing the values for the dimensions $\mathbf{d}_2$ of $\hat{x}$ on which $\mathcal{E}$ is mistaken.

A fair and adequate explanation must not only contain counterfactuals that resolve the explainee's conundrum. It must make clear the biases of the system which may account for $\mathcal{O}$'s incomplete understanding of $\hat{f}$; it must lay bare any prejudicial factors P that affect the explainee and thus in effect all overdetermining factors as in Definition 4. An explainee might reasonably want to know whether such biases resulted in a prediction concerning her. E.g., the explanation in (3) might satisfy CM or CI, but still be misleading. Thus:

CB $\forall$ prejudicial factors P , there is a $\Delta \in \mathfrak{C}_{\mathcal{E}}^d$ such that $\hat{f}(\Delta(\hat{x})) = \pi$ and $P(\Delta(\hat{x})) = \Delta(\hat{x})$.

In our bank loan example, if the bank is constrained to provide an explanation obeying CB, then it must provide an explanation according to which being white and having $\mathcal{E}$'s salary would have sufficed to get the loan.

Definition 11. *A set of counterfactuals provides a fair and adequate explanation of $\hat{f}$ for $\mathcal{E}$ at $\hat{x}$ just in case they together satisfy CM, CI and CB within a certain distance d of $\hat{x}$.*

The counterfactuals in $\mathfrak{C}_{\mathcal{E}}^d$ jointly provide a *fair and adequate* explanation of $\hat{f}$ for $\mathcal{E}$, though individually they may not satisfy all of the constraints. We investigate how hard it is to find an adequate local explanation in the next section.

5 The algorithmic complexity of finding fair and adequate explanations

In this section, we examine the computational complexity of finding a fair and adequate explanation. To find an appropriate explanation, we imagine a game played, say, between the bank and the would-be loan taker $\mathcal{E}$ in our example from Section 2, in which $\mathcal{E}$ can ask questions of the bank (or owner/ developer of the algorithm) about the algorithm's decisions. We propose to use a two player game, an *explanation game* to get appropriate explanations for the explainee.

The pragmatic nature of explanations already motivates the use of a game theoretic framework. We have argued fair and adequate explanations must obey pragmatic constraints; and in order to satisfy these in a cooperative game the explainer must understand explainee $\mathcal{E}$'s conundrum and respond so as to resolve it. Providing an explanation is a pragmatic act that takes into account an explainee's cognitive state and the conundrum it engenders for the particular fact that needs explaining. A cooperative explainer will provide an explanation in terms of the type he assigns to $\mathcal{E}$, as the type will encode the relevant portions of $\mathcal{E}$'s cognitive state. On the other hand the explainee will need to interpret the putative explanation in light of her model of the explainer's view of his type. Thus, both explainer and explainee naturally have strategies that exploit information about the other. Signaling games [37] are a well-understood and natural

formal framework in which to explore the interactions between explainer and explainee; the game theoretic machinery we develop below can be easily adapted into a signaling game between explainer and explainee where explanations succeed when their strategies coordinate on the same outcome.

Rather than develop signaling games however for coordinating on successful explanations, we look at non-cooperative scenarios where the explainer $\hat{f}$ may attempt to hide a good explanation. For instance, the bank in our running example might have encoded directly or indirectly biases into its loan program that are prejudicial to $\mathcal{E}$, and it might not want to expose these biases. The games below provide a formal account of the difficulty our explainee has in finding a winning strategy in such a setting.

To define an explanation game, we first fix a set of two players $\{\mathcal{E}, \mathcal{A}\}$.

The moves or actions $V_{\mathcal{E}}$ for explainee $\mathcal{E}$ are: playing an ACCEPT move—in which $\mathcal{E}$ accepts a proposed Δ_i if it partially solves her conundrum; playing an N-REQUEST move—i.e. requesting a Δ_j where j differs from all i such that Δ_i has been proposed by $\mathcal{A}$ in prior play; playing a P-REQUEST move—i.e. for some particular i , requesting Δ_i . $\mathcal{E}$ may also play a CHALLENGE move, in which $\mathcal{E}$ claims that a set of features $A_1, \dots, A_n$ of the focal point that entails π in the counterfactual model associated with $\hat{f}$. We distinguish three types of ME explanation games for $\mathcal{E}$ based on the types of moves she is allowed: the *Forcing* ME explanation games, in which $\mathcal{E}$ may play ACCEPT, N-REQUEST, P-REQUEST; the more restrictive *Restriction* ME explanation games, in which $\mathcal{E}$ may only play ACCEPT, N-REQUEST; and finally Challenge ME explanation games in which CHALLENGE moves are allowed.

Adversary $\mathcal{A}$'s moves $V_{\mathcal{A}}$ consists of the following: producing Δ_i and computing $\hat{f}(\Delta_i(\hat{x}))$ in response to N-REQUEST or P-REQUEST by $\mathcal{E}$; if $\mathcal{G}$ is a forcing game, $\mathcal{A}$ must play Δ_i at move m in ρ , if $\mathcal{E}$ has played P-REQUEST Δ_i at $m - 1$. In reacting to a N-REQUEST, player $\mathcal{A}$ may offer any new Δ_i ; if he is noncooperative, he will offer a new Δ_i that is not relevant to $\mathcal{E}$'s conundrum, unless he has no other choice. On the other hand, $\mathcal{A}$ must react to a CHALLENGE move by $\mathcal{E}$ by playing a Δ_i that either completes or corrects the Challenge assumption. A CHALLENGE demands a cooperative response; and since it can involve any implicitly definable prejudicial factor as in Definition 5, it can also establish CB, as well as remedy CI or CM.

We now specify a win-lose, generic *explanation game*.

Definition 12. An Explanation game, $\mathcal{G}$, concerning a polynomially computable function $\hat{f}: X^n \rightarrow Y$, where X^n is a space of boolean valued features for the data and Y a set of predictions, is a tuple $((V_{\mathcal{E}} \cup V_{\mathcal{A}})^*, \mathcal{E}, \mathcal{A}, \hat{f}: X^n \rightarrow Y, \hat{x}, d, \mathfrak{C}_{\mathcal{E}}^d)$ where:

- i. $\mathfrak{C}_{\mathcal{E}}^d \subseteq \mathbf{B}_{\hat{x}, \pi}^d$ resolves $\mathcal{E}$'s conundrum and obeys CB.
- ii. $\hat{x} \in X^n$ is the starting position, d is the antecedently fixed distance parameter.
- iii. $\mathcal{A}$, but not $\mathcal{E}$ has access to the behavior of $\hat{f}$ and a fortiori $\mathfrak{C}_{\mathcal{E}}^d$.
- iv. $\mathcal{E}$ opens $\mathcal{G}$ with a REQUEST or CHALLENGE move
- v. $\mathcal{A}$ responds to $\mathcal{E}$'s requests by playing some $\Delta_i, i \leq d$.

vi. $\mathcal{E}$ may either play *ACCEPT*, in which case the game ends or again play a *REQUEST* or *CHALLENGE* move.

$\mathcal{E}$ wins $\mathcal{G}$ just in case in $\mathcal{G}$ she can determine $\mathfrak{C}_{\mathcal{E}}^d$. The game terminates when (a) 0 has determined $\mathfrak{C}_{\mathcal{E}}^d$ (resolved her conundra) or gives up.

$\mathcal{E}$ always has a winning strategy in an explanation game. The real question is how quickly $\mathcal{E}$ can compute her winning condition. An answer depends on what moves we allow for $\mathcal{E}$ in the Explanation game; we can restrict $\mathcal{E}$ to playing a Restriction explanation game, a Forcing game or a Forcing game with CHALLENGE moves.

Proposition 5. *Suppose $\mathcal{G}$ is a forcing explanation game. Then the computation of $\mathcal{E}$'s winning strategy in $\mathcal{G}$ is Polynomial Local Search complete (PLS) [20,31]. On the other hand if $\mathcal{G}$ is only a Restriction game, then the worst case complexity for finding her strategy is exponential.*

Proof sketch. Finding $\mathfrak{C}_{\mathcal{E}}^d$ is a search problem using $\hat{f}$. $\mathfrak{C}_{\mathcal{E}}^d$ is finite with, say, m elements. These elements need not be unique; they just need jointly to solve the conundrum. This search problem is PLS just in case every solution element is polynomially bounded in the size of the input instance, $\hat{f}$ is poly-time, the cost of the solution is poly-time and it is possible to find the neighbors of any solution in poly-time. Let $\hat{x}$ be the input instance. By assumption, $\hat{f}$ is polynomial; and given the bound d , the solutions y for $\hat{f}(y) = \pi$ and $y \in \mathfrak{C}_{\mathcal{E}}^d$ are polynomially bounded in the size of the description of $\hat{x}$. Now, finding a point $y \in \mathfrak{C}_{\mathcal{E}}^d$ that solves at least part of $\mathcal{E}$'s conundrum, as well as finding neighbors of y is poly-time, since $\mathcal{E}$ can use P-REQUEST moves to direct the search. To determine the cost c of finding $\mathfrak{C}_{\mathcal{E}}^d$ for $|\mathfrak{C}_{\mathcal{E}}^d| = m$ in poly-time: we set for $y \in \mathfrak{C}_{\mathcal{E}}^d$ the j th element of $\mathfrak{C}$ computed as $c(y) = m - j$; if $y \notin \mathfrak{C}$, $c(y) = m$. Finding $\mathfrak{C}_{\mathcal{E}}^d$ thus involves determining m local minima and is PLS. In addition, determining $\mathfrak{C}_{\mathcal{E}}^d$ encodes the PLS complete problem FLIP [20]: the solutions y in $\mathcal{G}$ have the same edit distance as the solutions in FLIP, $\hat{f}$ encodes a starting position, and our cost function can be recoded over the values of the Boolean features defining y to encode the cost function of FLIP and the function that compares solutions in FLIP is also needed and constructible in $\mathcal{G}$. So finding $\mathfrak{C}_{\mathcal{E}}^d$ is PLS complete in $\mathcal{G}$ as it encodes FLIP.

The fact that forcing explanation games are PLS complete makes getting an appropriate explanation computationally difficult. Worse, if $\mathcal{G}$ is a Restriction Explanation game, then $\mathcal{A}$ can force $\mathcal{E}$ to enumerate all possible Δ_i within radius d of $\hat{x}$ to find $\mathfrak{C}_{\mathcal{E}}^d$. $\square$

Proposition 6. *Suppose $\mathcal{G}$ is a Challenge explanation game. Then $\mathcal{E}$ has a winning strategy in $\mathcal{G}$ that is linear time computable.*

Proof sketch. $\mathcal{A}$ must respond to $\mathcal{E}$'s CHALLENGE moves by correcting or completing $\mathcal{E}$'s proposed list of features. $\mathcal{E}$ can determine $\mathfrak{C}_{\mathcal{E}}^d$ in a number of moves that is linear in the size of $\mathfrak{C}_{\mathcal{E}}^d$. $\square$

A Challenge explanation game mimics a coordination game where $\mathcal{A}$ has perfect information about $\mathfrak{C}_{\mathcal{E}}^d$, because it forces cooperativity and coordination on the part of $\mathcal{A}$. Suppose $\mathcal{E}$ in our bank example claims that her salary should be sufficient for a loan. In response to the challenge, the bank could claim the salary is not sufficient; but that’s not true—the salary *is* sufficient *provided* other conditions hold. That is, $\mathcal{E}$ ’s conundrum is an instance of CI. Because of the constraint on CHALLENGE answers by the opponent, the bank must complete the missing element: *if you were white with a salary of €50K,...* Proposition 6 shows that when investigating an $\hat{f}$ in a challenge game, exploiting a conundrum is a highly efficient strategy.

The flip degree of $\mathbf{S}_{\hat{x},\pi}^d$ and the number of overdetermining factors $O(x, \pi)$ (Definition 4) typically affect the size of $\mathfrak{C}$ and thus the complexity of the conundrum and search for fair and adequate explanations and their logical valid associates. More particularly, when $|O(\pi, \hat{x})| = n$ and the cost of the prediction is as in the proof of Proposition 5, $\mathcal{E}$ ’s conundrum and the explanations resolving it may require n local minima. When the flip degree of $\mathbf{S}_{\hat{x},\pi}^d$ is m , $\mathcal{E}$ may need to compute m local minima.

To develop practical algorithms for fair and adequate explanations for AI systems, we need to isolate $\mathcal{E}$ ’s conundrum. This will enable us to exploit the efficiencies of Challenge explanation games. Extending the framework to discover $\mathcal{E}$ ’s conundrum behind her request for an explanation is something we plan to do using epistemic games from [3] with more developed linguistic moves. In a more restricted setting where Challenge games are not available, our game framework shows that clever search algorithms and heuristics for PLS problems will be essential to providing users with relevant, and provably fair and adequate counterfactual explanations. This is something current techniques like enumeration or finding closest counterparts, which may not be relevant [19,18,22]—do not do.

6 Conclusion

We have shown that counterfactual explanations can deliver partial, but epistemically accessible and adequate explanations. We have also shown that any counterfactual explanation can be extended to a valid deductive one. We have shown that pragmatic factors dramatically affect the complexity of finding adequate explanations, and we introduced Explanation Games, which provided to represent finding fair and adequate counterfactual explanations as a PLS complete search problem. In addition, we explored how the complexity of the set of counterfactuals describing a local neighborhood around the focal point can affect both the complexity of fair and adequate explanations and our evaluation of the learning algorithm as a model.

Our paper fills in part of the gap for finding fair and adequate explanations in a computationally reasonable way. Nevertheless moving from an explanation provided by an explanation game to a proof from a minimal set of sufficient premises as in Proposition 2 is still computationally difficult. In future work we will look at efficient heuristics for this step. In future work, we will also look at how explanation games help us to formally explore interactive machine learning, in

particular “human in the loop” or interactive explainability for machine learning function behavior [2,17]. Such game theoretic investigations may have special relevance in medical domains [16].

Acknowledgement

We thank the ANR PRCI grant SLANT, the ICT 38 EU grant COALA and the 3IA Institute ANITI funded by the ANR-19-PI3A-0004 grant for research support. We also thank the reviewers for their insightful comments.

References

1. Achinstein, P.: The Nature of Explanation. Oxford University Press (1980)
2. Amershi, S., Cakmak, M., Knox, W.B., Kulesza, T.: Power to the people: The role of humans in interactive machine learning. *Ai Magazine* **35**(4), 105–120 (2014)
3. Asher, N., Paul, S.: Strategic conversation under imperfect information: epistemic Message Exchange games. *Logic, Language and Information* **27.4**, 343–385 (2018)
4. Bachoc, F., Gamboa, F., Halford, M., Loubes, J.M., Risser, L.: Entropic variable projection for explainability and interpretability. *arXiv preprint arXiv:1810.07924* (2018)
5. Bromberger, S.: An approach to explanation. In: Butler, R. (ed.) *Analytical Philosophy*, pp. 72–105. Oxford University Press (1962)
6. Chang, C.C., Keisler, H.J.: *Model theory*. Elsevier (1990)
7. De Raedt, L., Dumančić, S., Manhaeve, R., Marra, G.: From statistical relational to neuro-symbolic artificial intelligence. *arXiv preprint arXiv:2003.08316* (2020)
8. Doshi-Velez, F., Kim, B.: Towards a rigorous science of interpretable machine learning. *arXiv preprint arXiv:1702.08608* (2017)
9. Dube, S.: High dimensional spaces, deep learning and adversarial examples. *arXiv preprint arXiv:1801.00634* (2018)
10. Fan, X., Toni, F.: On computing explanations in argumentation. In: Bonet, B., Koenig, S. (eds.) *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence*. pp. 1496–1502. AAAI Press (2015)
11. Friedrich, G., Zanker, M.: A taxonomy for generating explanations in recommender systems. *AI Magazine* **32**(3), 90–98 (2011)
12. Gärdenfors, P., Makinson, D.: Revisions of knowledge systems using epistemic entrenchment. In: Vardi, M.Y. (ed.) *Proceedings of the Second Conference on Theoretical Aspects of Reasoning about Knowledge*. pp. 83–95. Morgan Kaufmann, San Francisco (1988)
13. Ginsberg, M.L.: Counterfactuals. *Artificial intelligence* **30**(1), 35–79 (1986)
14. Hempel, C.G.: *Aspects of scientific explanation*. Free Press New York (1965)
15. Holzinger, A., Carrington, A., Müller, H.: Measuring the quality of explanations: the system causability scale (scs). *KI-Künstliche Intelligenz* pp. 1–6 (2020)
16. Holzinger, A., Malle, B., Saranti, A., Pfeifer, B.: Towards multi-modal causability with graph neural networks enabling information fusion for explainable ai. *Information Fusion* **71**, 28–37 (2021)
17. Holzinger, A., Plass, M., Kickmeier-Rust, M., Holzinger, K., Crişan, G.C., Pintea, C.M., Palade, V.: Interactive machine learning: experimental evidence for the human in the algorithmic loop. *Applied Intelligence* **49**(7), 2401–2414 (2019)

18. Ignatiev, A., Narodytska, N., Asher, N., Marques-Silva, J.: On relating “why?” and “why not?” explanations. In: *Proceedings of AI*IA 2020* (2020)
19. Ignatiev, A., Narodytska, N., Marques-Silva, J.: On relating explanations and adversarial examples. In: *Advances in Neural Information Processing Systems* (2019)
20. Johnson, D.S., Papadimitriou, C.H., Yannakakis, M.: How easy is local search? *Journal of computer and system sciences* **37**(1), 79–100 (1988)
21. Junker, U.: Preferred explanations and relaxations for over-constrained problems. In: *AAAI-2004* (2004)
22. Karimi, A.H., Barthe, G., Balle, B., Valera, I.: Model-agnostic counterfactual explanations for consequential decisions. In: *International Conference on Artificial Intelligence and Statistics*. pp. 895–905. PMLR (2020)
23. Kurakin, A., Goodfellow, I., Bengio, S.: Adversarial examples in the physical world. *arXiv preprint arXiv:1607.02533* (2016)
24. Kusner, M.J., Loftus, J., Russell, C., Silva, R.: Counterfactual fairness. In: *Advances in Neural Information Processing Systems*. pp. 4066–4076 (2017)
25. Laugel, T., Lesot, M.J., Marsala, C., Renard, X., Detyniecki, M.: Unjustified classification regions and counterfactual explanations in machine learning. In: *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*. pp. 37–54. Springer (2019)
26. Lewis, D.: Causation. *Journal of Philosophy* **70**(17), 556–567 (1973)
27. Lundberg, S.M., Lee, S.: A unified approach to interpreting model predictions. In: *NIPS*. pp. 4765–4774 (2017)
28. Miller, T.: Explanation in artificial intelligence: Insights from the social sciences. *Artificial Intelligence* pp. 1–38 (2019)
29. Molnar, C.: *Interpretable machine learning*. Lulu. com (2019)
30. Murdoch, W.J., Singh, C., Kumbier, K., Abbasi-Asl, R., Yu, B.: Definitions, methods, and applications in interpretable machine learning. *Proceedings of the National Academy of Sciences* **116**(44), 22071–22080 (2019)
31. Papadimitriou, C.H., Schäffer, A.A., Yannakakis, M.: On the complexity of local search. In: *Proceedings of the twenty-second annual ACM symposium on Theory of computing*. pp. 438–445 (1990)
32. Pearl, J.: System Z: a natural ordering of defaults with tractable applications to nonmonotonic reasoning. In: *Proceedings of the 3rd conference on Theoretical aspects of reasoning about knowledge (TARK’90)*. pp. 121–135 (1990)
33. Peyré, G., et al.: Computational optimal transport: With applications to data science. *Foundations and Trends in Machine Learning* **11**(5-6), 355–607 (2019)
34. Ribeiro, M.T., Singh, S., Guestrin, C.: ”why should I trust you?”: Explaining the predictions of any classifier. In: *KDD*. pp. 1135–1144 (2016)
35. Ribeiro, M.T., Singh, S., Guestrin, C.: Anchors: High-precision model-agnostic explanations. In: *AAAI*. pp. 1527–1535 (2018)
36. Salzberg, S.: Distance metrics for instance-based learning. In: *International Symposium on Methodologies for Intelligent Systems*. pp. 399–408. Springer (1991)
37. Spence, A.M.: Job market signaling. *Journal of Economics* **87**(3), 355–374 (1973)
38. Wachter, S., Mittelstadt, B., Russell, C.: Counterfactual explanations without opening the black box: Automated decisions and the gpdr. *Harv. JL & Tech.* **31**, 841 (2017)
39. Williamson, T.: First-order logics for comparative similarity. *Notre Dame Journal of Formal Logic* **29**(4) (1988)
40. Younes, L.: *Diffeomorphic learning* (2019), [arXiv:1806.01240](https://arxiv.org/abs/1806.01240)