A Proofs

A.1 Proof of Proposition 3.1

Proposition 3.1. *A CFE-based action, $\mathbf{A}^{\mathit{CFE}}$, where $I = \{i \mid \delta_i^* \neq 0\}$, performed by individual $\boldsymbol{x}^{\mathit{F}}$, in general results in the structural counterfactual, $\boldsymbol{x}^{\mathit{SCF}} = \boldsymbol{x}^{*\mathit{CFE}} := \boldsymbol{x}^{\mathit{F}} + \boldsymbol{\delta}^*$, and thus guarantees recourse (i.e., $h(\boldsymbol{x}^{\mathit{SCF}}) \neq h(\boldsymbol{x}^{\mathit{F}})$), if and only if, the set of descendants of the acted upon variables, determined by I, is the empty set.*

Proof. The setting assumes that the causal graph $\mathcal{G}$ is available such that the parent set for each variable is known. Let $\mathtt{d}(X)$ and $\mathtt{nd}(X)$ denote the sets of descendants and non-descendants of the variable X according to $\mathcal{G}$, respectively. For multiple intervened-upon variables, we define:

$$\mathbb{X}_I := \{X_i\}_{i \in I},$$
$$\mathtt{nd}(\mathbb{X}_I) := \cap_{i \in I} \mathtt{nd}(X_i),$$
$$\mathtt{d}(\mathbb{X}_I) := \mathbb{X} \setminus (\mathbb{X}_I \cup \mathtt{nd}(\mathbb{X}_I)).$$

Note that, by definition, $\mathbb{X}_I$, $\mathtt{nd}(\mathbb{X}_I)$, and $\mathtt{d}(\mathbb{X}_I)$ form a partition of the set of all variables $\mathbb{X}$.

To prove the iff conditional, we prove each direction separately. For ease of exposition, we define

$$\underbrace{\boldsymbol{x}^{\texttt{SCF}} = \boldsymbol{x}^{*\texttt{CFE}} := \boldsymbol{x}^{\texttt{F}} + \boldsymbol{\delta}^*}_{\mathbf{p}} \quad \Longleftrightarrow \quad \underbrace{\mathtt{d}(\mathbb{X}_I) = \varnothing}_{\mathbf{q}}$$

where we recall the remark that given $\boldsymbol{\delta}^*$, an individual seeking recourse may intervene on any arbitrary subset of observed variables $\mathbb{X}_I$, as long as $(\delta_i^* \neq 0) \implies (i \in I)$.

$\mathbf{q} \Longrightarrow \mathbf{p}$: Borrowing the closed-form expression of a structural counterfactual from (**??**), we have

$$x_i^{\texttt{SCF}} = \begin{cases} x_i^{\texttt{F}} + \delta_i^* & i \in I \\ x_i^{\texttt{F}} + f_i(\mathbf{pa}_i^{\texttt{SCF}}) - f_i(\mathbf{pa}_i^{\texttt{F}}) & i \notin I \end{cases} \tag{A.1}$$

which can be broken down further to specify the descendants and non-descendants of intervened upon variables, as

$$x_i^{\texttt{SCF}} = \begin{cases} x_i^{\texttt{F}} + \delta_i^* & i \in I \\ x_i^{\texttt{F}} + f_i(\mathbf{pa}_i^{\texttt{SCF}}) - f_i(\mathbf{pa}_i^{\texttt{F}}) & i \in \mathtt{d}(\mathbb{X}_I) \\ x_i^{\texttt{F}} + f_i(\mathbf{pa}_i^{\texttt{SCF}}) - f_i(\mathbf{pa}_i^{\texttt{F}}) & i \in \mathtt{nd}(\mathbb{X}_I) \end{cases} \tag{A.2}$$

By assumption, $\mathtt{d}(\mathbb{X}_I) = \varnothing$, so the second case never holds.

Furthermore, since structural interventions leave non-descendant variables unaffected, we have that

$$\mathbf{pa}_i^{\texttt{SCF}} = \mathbf{pa}_i^{\texttt{F}} \quad \forall i \in \mathtt{nd}(\mathbb{X}_I).$$

Consequently,

$$f_i(\mathbf{pa}_i^{\texttt{SCF}}) - f_i(\mathbf{pa}_i^{\texttt{F}}) = f_i(\mathbf{pa}_i^{\texttt{F}}) - f_i(\mathbf{pa}_i^{\texttt{F}}) = 0 \quad \forall i \in \mathtt{nd}(\mathbb{X}_I).$$

In summary, we have

$$x_i^{\texttt{SCF}} = \begin{cases} x_i^{\texttt{F}} + \delta_i^* & i \in I \\ x_i^{\texttt{F}} & i \in \mathtt{nd}(\mathbb{X}_I) \end{cases} \tag{A.3}$$

which, upon realising that $(\delta_i^* \neq 0) \implies (i \in I)$, reduces to $\boldsymbol{x}^{\texttt{SCF}} = \boldsymbol{x}^{*\texttt{CFE}} := \boldsymbol{x}^{\texttt{F}} + \boldsymbol{\delta}^*$ as desired.

$\neg\mathbf{q} \Longrightarrow \neg\mathbf{p}$: Starting with the negation of $\mathbf{q}$, we have the $\exists\, k \in I$ s.t. $\mathtt{d}(X_k) \neq \varnothing$. It is assumed that $\delta_k^* \neq 0$ (i.e., we are not performing a non-altering intervention on X_k), then using the same expression for structural counterfactuals in (A.2), there in general exists a descendant of X_k for which the value of its ancestors change under intervention, i.e., $\exists\, l \in \mathtt{d}(\mathbb{X}_I)$ s.t. $f_l(\mathbf{pa}_l^{\texttt{SCF}}) - f_l(\mathbf{pa}_l^{\texttt{F}}) \neq 0$. Thus, $x_l^{\texttt{SCF}} \neq x_l^{\texttt{F}}$ and thus $\boldsymbol{x}^{\texttt{SCF}} \neq \boldsymbol{x}^{*\texttt{CFE}} := \boldsymbol{x}^{\texttt{F}} + \boldsymbol{\delta}^*$. Our proof ignores special cases such as piece-wise constant structural equations, where for some $\delta_i^* \neq 0$, the descendant of X_i remains invariant. These rare cases can be thought of as locally violating causal minimality [36, Sec. 6.5] and are thus disregarded. □

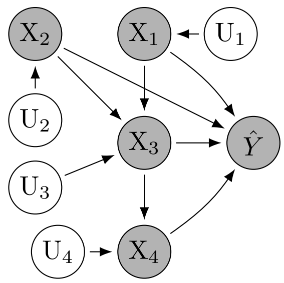

$$\left.\begin{aligned} X_1 &:= U_1 \\ X_2 &:= U_2 \\ X_3 &:= f_3(X_1, X_2) + U_3 \\ X_4 &:= f_4(X_3) + U_4 \end{aligned}\right\} \mathcal{M}$$

$$\hat{Y} = h\big(\{X_i\}_{i=1}^4\big)$$

Figure 3: The structural causal model (graph and equations) for the working example and demonstration in Section 4.

A.2 Proof of Corollary 3.1

Corollary 3.1. *If the true world $\mathcal{M}$ is independent, i.e, all the observed features are root-nodes, then CFE-based actions always guarantee recourse.*

Proof. If the true world $\mathcal{M}$ is independent, then by definition the set of descendants for all variables is the empty set. Thus, the statement follows directly from Proposition 3.1. □

A.3 Proof of Proposition 4.1

Proposition 4.1. *Given an individual $\boldsymbol{x}^{\texttt{F}}$ observed in world $\mathcal{M} \in \Pi$, a family of feasible actions $\mathcal{F}$, and the solution of* (3), $\mathbf{A}^* \in \mathcal{F}$. *Assume that there exists CFE-based action $\mathbf{A}^{\texttt{CFE}} \in \mathcal{F}$ that achieves recourse, i.e., $h(\boldsymbol{x}^{\texttt{F}}) \neq h(\boldsymbol{x}^{*\texttt{CFE}})$. Then,* $\mathrm{cost}(\mathbf{A}^*; \boldsymbol{x}^{\texttt{F}}) \leq \mathrm{cost}(\mathbf{A}^{\texttt{CFE}}; \boldsymbol{x}^{\texttt{F}})$.

Proof. Having assumed that both $\mathbf{A}^{\texttt{CFE}}, \mathbf{A}^* \in \mathcal{F}$, and considering that $\mathbf{A}^*$ is the optimal solution of (3) constrained to $\mathcal{F}$, it follows from definition of optimality that $\mathrm{cost}(\mathbf{A}^*; \boldsymbol{x}^{\texttt{F}}) \leq \mathrm{cost}(\mathbf{A}^{\texttt{CFE}}; \boldsymbol{x}^{\texttt{F}})$. □

B Working example

Consider the model in Figure 3, where $\{U_i\}_{i=1}^4$ are mutually independent exogenous variables, and $\{f_i\}_{i=1}^4$ are structural (linear or nonlinear) equations. Let $\boldsymbol{x}^{\texttt{F}} = [x_1^{\texttt{F}}, x_2^{\texttt{F}}, x_3^{\texttt{F}}, x_4^{\texttt{F}}]^T$ be the observed features belonging to the (factual) individual, for whom we seek a counterfactual explanation and recommendation. Also, let I denote the set of indices corresponding to the subset of endogenous variables that are intervened upon according to the action set $\mathbf{A}$. Then, we obtain a structural counterfactual, $\boldsymbol{x}^{\texttt{SCF}} = \mathbb{F}_{\mathbf{A}}(\mathbb{F}^{-1}(\boldsymbol{x}^{\texttt{F}}))$, by applying the Abduction-Action-Prediction steps [34] as follows:

Step 1. Abduction uniquely determines the value of all exogenous variables, $\{u_i\}_{i=1}^4$, given evidence, $\{X_i = x_i^{\texttt{F}}\}_{i=1}^4$:

$$\begin{aligned} u_1 &= x_1^{\texttt{F}}, \\ u_2 &= x_2^{\texttt{F}}, \\ u_3 &= x_3^{\texttt{F}} - f_3(x_1^{\texttt{F}}, x_2^{\texttt{F}}), \\ u_4 &= x_4^{\texttt{F}} - f_4(x_3^{\texttt{F}}). \end{aligned} \tag{B.1}$$

Step 2. Action modifies the SCM according to the hypothetical interventions, $\mathrm{do}(\{X_i := a_i\}_{i \in I})$ (where $a_i = x_i^F + \delta_i$), yielding $\mathbb{F}_{\mathbf{A}}$:

$$\begin{aligned} X_1 &:= [1 \in I] \cdot a_1 + [1 \notin I] \cdot U_1, \\ X_2 &:= [2 \in I] \cdot a_2 + [2 \notin I] \cdot U_2, \\ X_3 &:= [3 \in I] \cdot a_3 + [3 \notin I] \cdot \big(f_3(X_1, X_2) + U_3\big), \\ X_4 &:= [4 \in I] \cdot a_4 + [4 \notin I] \cdot \big(f_4(X_3) + U_4\big), \end{aligned} \tag{B.2}$$

where $[\cdot]$ denotes the Iverson bracket.

Step 3. Prediction recursively determines the values of all endogenous variables based on the computed exogenous variables $\{u_i\}_{i=1}^4$ from Step 1 and $\mathbb{F}_{\mathbf{A}}$ from Step 2, as:

$$\begin{aligned}
x_1^{\texttt{SCF}} &:= [1 \in I] \cdot a_1 + [1 \notin I] \cdot \big(u_1\big), \\
x_2^{\texttt{SCF}} &:= [2 \in I] \cdot a_2 + [2 \notin I] \cdot \big(u_2\big), \\
x_3^{\texttt{SCF}} &:= [3 \in I] \cdot a_3 + [3 \notin I] \cdot \big(f_3(x_1^{\texttt{SCF}}, x_2^{\texttt{SCF}}) + u_3\big), \\
x_4^{\texttt{SCF}} &:= [4 \in I] \cdot a_4 + [4 \notin I] \cdot \big(f_4(x_3^{\texttt{SCF}}) + u_4\big).
\end{aligned} \tag{B.3}$$

C Demonstration

We showcase our proposed formulation by comparing the actions recommended by existing (nearest) counterfactual explanation methods, as in (2), to the ones generated by the proposed minimal intervention formulation in (3). We recall that prior literature has focused on generating counterfactual explanations or CFE-based actions, which as shown above lack optimally or feasibility guarantees in non-independent worlds. Thus, to the best of our knowledge, there exists no baseline approach in the literature that guarantees algorithmic recourse. The experiments below serve as an illustration of the sub-optimality of existing approaches relative to our proposed formulation of recourse via minimal intervention. Section 5 presents a detailed discussion on practical considerations.

We consider two settings: i) a synthetic setting where $\mathcal{M}$ follows Figure 1; and ii) a real-world setting based on the `german` credit dataset [1], where $\mathcal{M}$ follows Figure 3. We computed the cost of actions as the ℓ_1 norm over normalized feature changes to make effort comparable across features, i.e., $\text{cost}(\cdot; \boldsymbol{x}^{\texttt{F}}) = \sum_{i \in I} |\delta_i|/R_i$, where R_i is the range of feature i.

For the *synthetic setting*, we generate data following the model in Figure 1, where we assume $\mathrm{X}_1 := \mathrm{U}_1$, $\mathrm{X}_2 := 3/10 \cdot \mathrm{X}_1 + \mathrm{U}_2$, with $\mathrm{U}_1 \sim \$10000 \cdot \text{Poission}(10)$ and $\mathrm{U}_2 \sim \$2500 \cdot \mathcal{N}(0,1)$; and the predictive model $h = \text{sgn}(\mathrm{X}_1 + 5 \cdot \mathrm{X}_2 - \$225000)$. Given $\boldsymbol{x}^{\texttt{F}} = [\$75000, \$25000]^T$, solving our formulation, (3), identifies the optimal action set $\mathbf{A}^* = \text{do}(\mathrm{X}_1 := x_1^{\texttt{F}} + \$10000)$ which results in $\boldsymbol{x}^{*\texttt{SCF}} = \mathbb{F}_{\mathbf{A}^*}(\mathbb{F}^{-1}(\boldsymbol{x}^{\texttt{F}})) = [\$85000, \$28000]^T$, whereas solving previous formulations, (2), yields $\boldsymbol{\delta}^* = [\$0, +\$5000]^T$ resulting in $\boldsymbol{x}^{*\texttt{CFE}} = \boldsymbol{x}^{\texttt{F}} + \boldsymbol{\delta}^* = [\$75000, \$30000]^T$. Importantly, while $\boldsymbol{x}^{*\texttt{SCF}}$ appears to be at a further distance from $\boldsymbol{x}^{\texttt{F}}$ compared to $\boldsymbol{x}^{*\texttt{CFE}}$, achieving the former is less costly than the latter, specifically, $\text{cost}(\boldsymbol{\delta}^*; \boldsymbol{x}^{\texttt{F}}) \approx 2\, \text{cost}(\mathbf{A}^*; \boldsymbol{x}^{\texttt{F}})$.

As a *real-world setting*, we consider a subset of the features in the `german` credit dataset. The setup is depicted in Figure 3, where X_1 is the individual's gender (treated as immutable), X_2 is the individual's age (actionable but can only increase), X_3 is credit given by the bank (actionable), X_4 is the repayment duration of the credit (non-actionable but mutable), and $\hat{\mathrm{Y}}$ is the predicted customer risk, according to h (logisitic regression or decision tree). We learn the structural equations by fitting a linear regression model to the child-parent tuples. We will release the data, and the code used to learn models and structural equations.

Given the setup above, for instance, for the individual $\boldsymbol{x}^{\texttt{F}} = [\texttt{Male}, 32, \$1938, 24]^T$ identified as a risky customer, solving our formulation, (3), yields the optimal action set $\mathbf{A}^* = \text{do}(\{\mathrm{X}_2 := x_2^{\texttt{F}} + 1, \mathrm{X}_3 := x_3^{\texttt{F}} - \$800\})$ which results in $\boldsymbol{x}^{*\texttt{SCF}} = \mathbb{F}_{\mathbf{A}^*}(\mathbb{F}^{-1}(\boldsymbol{x}^{\texttt{F}})) = [\texttt{Male}, 33, \$1138, 22]^T$, whereas solving (2) yields $\boldsymbol{\delta}^* = [\texttt{N/A}, +6, 0, 0]^T$ resulting in $\boldsymbol{x}^{*\texttt{CFE}} = \boldsymbol{x}^{\texttt{F}} + \boldsymbol{\delta}^* = [\texttt{Male}, 38, \$1938, 24]^T$. Similar to the toy setting, we observe a $\%42$ decrease in effort required of the individual when using the action by our method, since our cost function states that waiting for six years to get the credit approved is more costly than applying the following year for a lower ($-\$800$) credit amount. We extend our analysis to a population level, and observe that for 50 negatively affected test individuals, previous approaches suggest actions that are on average $\%39 \pm \%24$ and $\%65 \pm \%8$ more costly than our approach when considering, respectively, a logistic regression and a decision tree as the predictive model h.

The demonstrations above confirm our theoretical analysis that MINT-based actions from (3) are less costly and thus more beneficial for affected individuals than existing CFE-based actions from (2) that fail to utilize the causal relations between variables.

D Towards Realistic Interventions

In Section 4, we formulated algorithmic recourse by considering the causal relations between features in the real world. Our formulation minimized the cost of actions, which were carried out as *structural* interventions on the corresponding graph. Each intervention proceeds by *unconditionally severing all edges* incident on the intervened node, fixing the post-manipulation distribution of a *single* variable to *one deterministic* value. While intuitive appealing and powerful, structural interventions are in many ways the simplest type of interventions, and their "simplicity comes at a price: foregoing the possibility of modeling many situations realisitically" [8, 22]. Below, we extend (3) and (4) to add flexibility and realism to the types of interventions performed by the individual. Notably, there is nothing inherent to an SCM that a priori determines the *form*, *feasibility*, or *scope* of intervention; instead, these choices are delegated to the individual and are made based on a semantic understanding of the modeled variables.

D.1 On the Form of Interventions

The demonstrations in Section C primarily focused on actions performed as *structural (a.k.a., hard)* interventions [33] where all incoming edges to the intervened node are severed (see (4)). Hard interventions are particularly useful for Randomized Control Trial (RCT) settings where one aims to evaluate (isolate) the causal effect of an action (e.g., effect of aspirin on patients with migraine) on the population by randomly assigning individuals to treatment/control groups, removing the influence of other factors (e.g., age).

In the context of algorithmic recourse, however, an individual performs actions in the real world, and therefore must play the rules governing the world. In earlier sections, these rules (captured in an SCM) guided the search for an optimal set of actions by modelling actions along with their consequences. The rules also determine the form of an intervention, e.g., specifying whether an intervention cancels out or complements existing causal relations.

For instance, consider **Example 1**, where an individual chooses to increase their bank balance (e.g., through borrowing money from family, i.e., a deliberate action/intervention on X_2 while continuing to put aside a portion of their income (i.e., retaining the relation $\mathrm{X}_2 := 3/10 \cdot \mathrm{X}_1 + \mathrm{U}_2$). Indeed, it would be unwise for a recommendation to suggest abandoning saving habits. In such a scenario, the action would be carried out as an *additive (a.k.a., soft)* intervention [10]. Such interventions *do not* sever graphical edges incident on the intervened node and continue to allow for parents of the node to affect that node. Conversely, in **Example 2**, recourse recommendations may suggest performing a structural intervention on temperature, e.g., by creating a climate controlled green-house, to cancel the natural effect of altitude change on temperature.

The previous examples illustrate a scenario where an individual/agriculture team actually have the agency to choose which type of intervention to perform. However, it is easy to conceive of examples where such an option does not exist. For instance, as part of a medical system's recommendation, we might consider adding $5\ \mathrm{mg/l}$ of insulin to a patient with diabetes with a certain blood insulin level [35]. This action cannot disable pre-existing mechanisms regulating blood insulin levels and therefore, the action can only be performed additively. Conversely, one may also consider another example from the medical domain whereby the only treatment of malignancy may be through a surgical (structural) amputation.[6]

Just as structural interventions were supported in our framework via a closed-form expression (see (4)), additive interventions can be encoded through an analogous assignment formulation:

$$x_i^{\mathtt{SCF}} = [i \in I] \cdot \delta_i + \big(x_i^{\mathtt{F}} + f_i(\mathbf{pa}_i^{\mathtt{SCF}}) - f_i(\mathbf{pa}_i^{\mathtt{F}})\big). \tag{D.1}$$

The choice of whether interventions should be applied in a additive/soft or structural/hard manner depends on the variable semantic [3], and should be decided prior to solving (3).

D.2 On the Feasibility of Interventions

We saw in Section 3 that earlier works motivated the addition of *feasibility* constraints as a means to provide more actionable recommendations for the individual seeking recourse [48]. There, the

[6] See, e.g., `https://www.cancer.org/cancer/bone-cancer/treating/surgery.html`.

actionability (a.k.a. *mutability*) of a feature was determined based on the feature semantic and value in the factual instance, marking those features which the individual has/lacks the agency to change (e.g., bank balance vs. race). While the interchangeable use of definition holds under an independent world, it fails when operating in most real-world settings governed by a set of causal dependencies. We study this subtlety below.

In an independent world, any change to variable X_i could come about only via an intervention on X_i itself. Therefore, immutable and non-actionable variables overlap. In a dependent world, however, changes to variable X_i may arise from an intervention on X_i or through changes to any of the ancestors of X_i. In this more general setting, we can tease apart the definition of *actionability* and *mutability*, and distinguish between three types of variables: (i) immutable (and hence non-actionable), e.g., race; (ii) mutable but non-actionable, e.g., credit score; and (iii) actionable (and hence mutable), e.g., bank balance. Each type requires special consideration which we show can be intuitively encoded as constraints amended to $\mathbf{A} \in \mathcal{F}$ from (3).

Immutable: We posit that the set of immutable (and hence non-actionable) variables should be closed under ancestral relationships given by the model, $\mathcal{M}$. This condition parallels the ancestral closure of *protected* attributions in [23]. This would ensure that under no circumstance would an intervention on an ancestor of an immutable variable change the immutable variable. Therefore, for an immutable variable X_i, the constraint $[i \notin I] = 1$ recursively necessitates the fulfillment of additional constraints $[j \notin I] = 1 \ \forall \ j \in \mathbf{pa}_i$ in $\mathcal{F}$. For instance, the immutability of race triggers the immutability of birthplace.

Mutable but non-actionable: To encode the conditions for mutable but non-actionable variables, we note that while a variable may not be directly actionable, it may still change as a result of changes to its parents. For example, the financial credit score in Figure 3 may change as a result of interventions to salary or savings, but is not itself directly intervenable. Therefore, for a non-actionable but mutable variable X_i, the constraint $[i \notin I] = 1$ is sufficient and does not induce any other constraints.

Actionable: In the most general sense, the actionable feasibility of an intervention on X_i may be contingent on a number of conditions, as follows: (a) the pre-intervention value of the intervened variable (i.e., $x_i^{\texttt{F}}$); (b) the pre-intervention value of other variables (i.e., $\{x_j^{\texttt{F}}\}_{j \subset [d] \setminus i}$); (c) the post-intervention value of the intervened variable (i.e., $x_i^{\texttt{SCF}}$); and (d) the post-intervention value of other variables (i.e., $\{x_j^{\texttt{SCF}}\}_{j \subset [d] \setminus i}$). Such feasibility conditions can easily be encoded into $\mathcal{F}$; consider the following scenarios:

(a) an individual's age can only increase, i.e., $[x_{age}^{\texttt{SCF}} \geq x_{age}^{\texttt{F}}]$; (b) an individual cannot apply for credit on a temporary visa, i.e.,
$[x_{visa}^{\texttt{F}} = \texttt{PERMANENT}] \geq [x_{credit}^{\texttt{SCF}} = \texttt{TRUE}]$;

(c) an individual may undergo heart surgery (an additive intervention) only if they won't remiss due to sustained smoking habits, i.e., $[x_{heart}^{\texttt{SCF}} \neq \texttt{REMISSION}]$; and

(d) an individual may undergo heart surgery only *after* their blood pressure is regularized due to medicinal intervention, i.e., $[x_{bp}^{\texttt{SCF}} = \texttt{O.K.}] \geq [x_{heart}^{\texttt{SCF}} = \texttt{SURGERY}]$.

In summary, while previous works on algorithmic recourse distinguished between actionable, conditionally actionable,[7] and immutable variables [48], we can now operate on a more realistic *spectrum* of variables, ranging from conditionally soft/hard actionable, to non-actionable but mutable, and finally to immutable and non-actionable variables. Finally, we remind that feasibility is a distinct notion from plausibility; whereas the former restricts actions $\mathbf{A} \in \mathcal{F}$ to those that can be performed by the individual, the latter determines the likeliness of the counterfactual instance $\boldsymbol{x}^{\texttt{SCF}} = \mathbb{F}_{\mathbf{A}}(\mathbb{F}^{-1}(\boldsymbol{x}^{\texttt{F}})) \in \mathcal{P}$ resulting from those actions. For instance, building on the earlier example, although an individual with similar attributes and higher credit score may exist in the dataset (i.e., plausible), directly acting on credit score is not feasible.

[7] Ustun et al. [48] also support conditionally actionable features (e.g., age or educational degree) with conditions derived only from $x_i^{\texttt{F}}$ as in (a). We generalize the set of conditions to support actions conditioned on the value of other variables as in (b), additive interventions in (c), and sequential interventions as in (d).

D.3 On the Scope of Interventions

One final assumption has been made throughout our discussion of actions as interventions which pertain to the one-to-one mapping between an action in the real world and an intervention on a endogenous variable in the structural causal model (which in turn are also input features to the predictive model). As exemplified in [3], it is possible for some actions (e.g., finding a higher-paying job) to simultaneously intervene on multiple variables in the model (e.g., income and length of employment). Alternatively, for **Example 2**, choosing a new paddy location is equivalent to intervening jointly on several input features of the predictive model (e.g., altitude, radiation, precipitation). Such confounded/correlated interventions, referred to as *fat-hand/non-atomic* interventions [10], will be explored further in follow-up work, by modelling the world at different causally consistent levels [4, 39].