



# Synthesizing explainable counterfactual policies for algorithmic recourse with program synthesis

Giovanni De Toni<sup>1,2</sup> · Bruno Lepri<sup>1</sup> · Andrea Passerini<sup>2</sup>

Received: 1 May 2022 / Revised: 21 November 2022 / Accepted: 7 December 2022 /  
Published online: 2 February 2023  
© The Author(s) 2023

## Abstract

Being able to provide counterfactual interventions—sequences of actions we would have had to take for a desirable outcome to happen—is essential to explain how to change an unfavourable decision by a black-box machine learning model (e.g., being denied a loan request). Existing solutions have mainly focused on generating feasible interventions without providing explanations of their rationale. Moreover, they need to solve a separate optimization problem for each user. In this paper, we take a different approach and learn a program that outputs a sequence of explainable counterfactual actions given a user description and a causal graph. We leverage program synthesis techniques, reinforcement learning coupled with Monte Carlo Tree Search for efficient exploration, and rule learning to extract explanations for each recommended action. An experimental evaluation on synthetic and real-world datasets shows how our approach, FARE (eFficient counterfActual REcourse), generates effective interventions by making orders of magnitude fewer queries to the black-box classifier with respect to existing solutions, with the additional benefit of complementing them with interpretable explanations.

**Keywords** Algorithmic recourse · Counterfactuals examples · Explainable AI · Machine learning

---

Editors: Alireza Tamaddon-Nezhad, Alan Bundy, Luc De Raedt, Artur d'Avila Garcez, Sebastijan Dumančić, César Ferri, Pascal Hitzler, Nikos Katzouris, Denis Mareschal, Stephen Muggleton, Ute Schmid.

---

✉ Giovanni De Toni  
giovanni.detoni@unitn.it

Bruno Lepri  
lepri@fbk.eu

Andrea Passerini  
andrea.passerini@unitn.it

<sup>1</sup> Fondazione Bruno Kessler, Via Sommarive 18, 38123 Trento, Italy

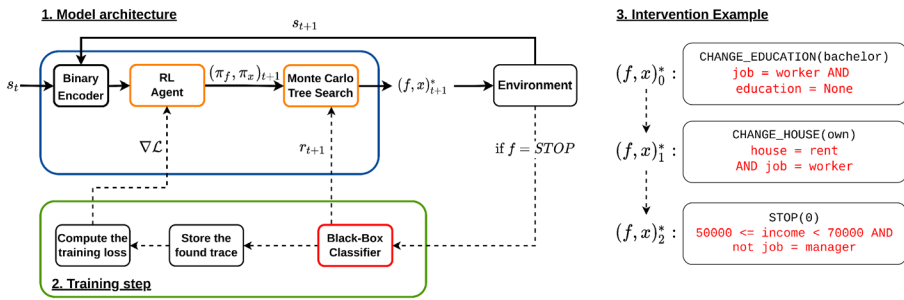
<sup>2</sup> Department of Information Engineering and Computer Science, University of Trento, Via Sommarive 9, 38123 Trento, Italy

## 1 Introduction

Counterfactual explanations are very powerful tools to explain the decision process of machine learning models (Wachter et al., 2017; Karimi et al., 2020). They give us the intuition of what could have happened if the state of the world was different (e.g., if you had taken the umbrella, you would not have gotten soaked). Researchers have developed many methods that can generate counterfactual explanations given a trained model (Wachter et al., 2017; Dandl et al., 2020; Mothilal et al., 2020; Karimi et al., 2020; Guidotti et al., 2018; Stepin et al., 2021). However, these methods do not provide any actionable information about which steps are required to obtain the given counterfactual. Thus, most of these methods do not enable algorithmic recourse. Algorithmic recourse describes the ability to provide “explanations and recommendations to individuals who are unfavourably treated by automated decision-making systems” (Karimi et al., 2021). For instance, algorithmic recourse can answer questions such as: what actions does a user have to perform to be granted a loan? Recently, providing feasible algorithmic recourse has also become a legal necessity (Voigt & Bussche, 2017). Some research works address this problem by developing ways to generate counterfactual interventions (Karimi et al., 2021), i.e., sequences of actions that, if followed, can overturn a decision made by a machine learning model, thus guaranteeing recourse. While being quite successful, these methods have several limitations. First, they are purely optimization methods that must be rerun from scratch for each new user. As a consequence, this requirement prevents their use for real-time intervention generation. Second, they are expensive in terms of queries to the black-box classifier and computing time. Last but not least, they fail to explain their recommendations (e.g., why does the model suggest getting a better degree rather than changing jobs?). On the contrary, explainability has been pointed out as a major requirement for methods generating counterfactual interventions (Barocas et al., 2020).

In this paper, we cast the problem of providing explainable counterfactual interventions as a program synthesis task (De Toni et al., 2021; Pierrot et al., 2019; Bunel et al., 2018; Balog et al., 2017): we want to generate a “program” that provides all the steps needed to overturn a bad decision made by a machine learning model. We propose a novel reinforcement learning (RL) method coupled with a discrete search procedure, Monte Carlo Tree Search (Coulom, 2006), to generate counterfactual interventions in an efficient data-driven manner. We call it FARE (eFficient counterfActual REcourse). As done by Naumann and Ntoutsis (2021), we assume a causal model encoding relationships between user features and consequences of potential interventions. We also provide a solution to distil an explainable deterministic program from the learned policy in the form of an automaton (E-FARE, Explainable and eFficient counterfActual REcourse). Figure 1 provides an overview of the architecture and the learning strategy and an example of an explainable intervention generated by the extracted automaton. Our approach addresses the three main limitations characterizing existing solutions:

- It learns a general policy that can be used to generate interventions for multiple users, rather than running separate user-specific optimizations.
- By coupling reinforcement learning with Monte Carlo Tree Search, it can efficiently explore the search space, requiring massively fewer queries to the black-box classifier than the best evolutionary algorithm (EA) model available, especially in settings with many features and (relatively) long interventions.



**Fig. 1** 1. Model architecture. Given the state  $s_t$  representing the features of the user, the agent generates candidate intervention policies  $\pi_f$  and  $\pi_x$  for functions and arguments, respectively (an action is a function-argument pair). MCTS uses these policies as a prior, and it extracts the best next action  $(f, x)_{t+1}^*$ . Once found, the reward received upon making the action is used to improve the MCTS estimates, and correct traces (i.e., those leading to the desired outcome change) are saved in a replay buffer. 2. Training step. The buffer is used to sample a subset of correct traces to be used to train the RL agent to mimic the behaviour of MCTS. 3. Explainable intervention. Example of an explainable intervention generated by the automaton extracted from the learned agent. Actions are in black, while explanations for each action are in red

- By extracting a program from the learned policy, it can complement the intervention with explanations motivating each action from contextual information. Furthermore, the program can be executed in real-time without accessing the black-box classifier.

Our experimental results on synthetic and real-world datasets confirm the advantages of the proposed solution over existing alternatives in terms of generality, scalability and interpretability.

## 2 Related work


Counterfactual explanations are versatile techniques to provide post-hoc interpretability of black-box machine learning models (Wachter et al., 2017; Dandl et al., 2020; Mothilal et al., 2020; Karimi et al., 2020; Guidotti et al., 2018; Stepin et al., 2021). They are model-agnostic, which means that they can be applied to trained models without performance loss. Compared to other global methods (Greenwell et al., 2018; Apley & Zhu, 2020), they provide instead *local* explanations. Namely, they underline only the *relevant factors* impacting a decision for a given initial target instance. They are also human-friendly and present many characteristics of what it is considered to be a *good explanation* (Miller, 2019). Therefore, they are suitable candidates to provide explanations to end-users since they are both highly-informative and localized. Recent research has shown how to generate counterfactual interventions for algorithmic recourse via various techniques (Karimi et al., 2020), such as probabilistic models (Karimi et al., 2020), integer programming (Ustun et al., 2019; Kanamori et al., 2020), reinforcement learning (Yonadav & Moses, 2019), program synthesis (Ramakrishnan et al., 2020), and genetic algorithms (Naumann & Ntoutsis, 2021). Researchers also developed solutions tied to a specific class of machine learning models, such as linear models (Tolomei et al., 2017) or Additive Tree Models (Cui et al., 2015). Methods with (approximated) convergence guarantees on the optimal counterfactual policies have also been proposed (Tsirtsis & Rodriguez, 2020). However, most of these methods ignore the causal relationships between user features (Tsirtsis &

Rodriguez, 2020; Ustun et al., 2019; Yonadav & Moses, 2019; Ramakrishnan et al., 2020). Without assuming an underlying causal graph, the proposed interventions become *permutation invariant*. For example, given an intervention consisting of three actions  $[A, B, C]$ , any intervention that is a permutation of the actions will have the same total cost. More importantly, it has been recently shown that optimal algorithmic recourse is impossible to achieve without a causal model of the interactions between the features (Karimi et al., 2020). The work by Karimi et al. (2020) provides algorithmic recourse following a probabilistic causal model but optimizes for subpopulation-based interventions instead of personalizing for a single user. CSCF (Naumann & Ntoutsis, 2021) is the only model-agnostic method capable of producing *consequence-aware* sequential interventions by exploiting causal relationships between features represented by a causal graph. However, CSCF is still purely an (evolutionary-based) optimization method, so it has to be run from scratch for each new user. Furthermore, the approach is opaque with respect to the reasons behind a suggested intervention. In this work, we show how our approach improves over CSCF in terms of generality, efficiency and interpretability.

## 3 Methods

### 3.1 Problem setting

The state of a user is represented as a vector of attributes  $s \in \mathcal{S}$  (e.g., age, sex, monthly income, job). A black-box classifier  $h : \mathcal{S} \rightarrow \{True, False\}$  predicts an outcome given a user state, with *True* being favourable to the user and *False* being unfavourable. The setting can be easily extended to multiclass classification by either grouping outcomes in favourable and unfavourable ones or learning separate programs converting from one class to the other. A *counterfactual intervention*  $I$  is a sequence of actions. Each action is represented as a tuple,  $(f, x) \in \mathcal{A}$ , composed by a *function*,  $f$ , and its argument,  $x \in \mathcal{X}_f$  (e.g.,  $(change\_income, 500)$ ). When an action is performed for a certain user, it modifies their state by altering one of their attributes according to its argument. A library  $\mathcal{F}$  contains all the possible functions which can be called. This library and the corresponding DSL (Domain Specific Language) are typically defined as a-priori by experts to prevent changes to protected attributes (e.g., age, sex, etc.). Examples of such DSLs can be found in "Appendix B". Moreover, each function possesses *pre-conditions* in the form of Boolean predicates over its arguments which describe the conditions that a user state must meet in order for a function to be called. The end of an intervention  $I$  is always specified by the `STOP` action. We also define a cost function,  $C : \mathcal{A} \times \mathcal{S} \rightarrow \mathbb{R}$  which mimics the effort made by a given user to perform an action given the current state. The cost is computed by looking at a causal graph  $\mathcal{G}$  (Pearl, 2009), where the nodes of the graph are the user's features. This assumption encodes the concept of *consequences* and it ensures a notion of order for the intervention's actions. For example, it might be easier to get first a degree and then a better salary rather than doing the opposite. The causal graph is problem-specific, and we can estimate it using domain knowledge or a domain expert. If we have observational data, we can also try to learn a candidate  $\mathcal{G}$  using automated methods (Tian & Pearl, 2001; Spirtes & Zhang, 2016), although inferring the "true" causal graph without interventions is not trivial. We use the former method for the evaluation by manually crafting the causal graphs. Figure 2 shows an example of a causal graph  $\mathcal{G}$  and of the corresponding costs. Our goal is to train an agent that, given a user with an unfavourable outcome, generates counterfactual



**Fig. 2** Examples of interventions on a causal graph. **A** A causal graph and a set of candidate actions. **B** Examples of interventions together with their costs. Note that the green line ( $\sum C = 15$ ) has a lower cost than the red line ( $\sum C = 28$ ) thanks to a better ordering of the actions making up the intervention (Color figure online)

interventions that overturn it. Given a black-box classifier  $h$ , a user  $s_0$  for whom the prediction by  $h$  is unfavourable (i.e.,  $h(s_0) = \text{False}$ ), a causal graph  $\mathcal{G}$  and a set of possible actions  $\mathcal{A}$  (implicitly represented by the functions in  $\mathcal{F}$  and their arguments in  $\mathcal{X}$ ), we want to generate a sequence  $I^*$ , that, if applied to  $s_0$ , produces a new state,  $s^* = I(s_0)$ , such that  $h(s^*) = \text{True}$ . This sequence must be *actionable*, which means that the user has to be able to perform those actions, and minimize the user’s *cost*. More formally:


$$\begin{aligned}
 I^* &= \min_I \sum_{t=0}^T C(a_t, s_t) \\
 \text{s.t.} \quad &I = \{a_t\}_{t=0}^T \quad a_t \in \mathcal{A} \quad \forall t \\
 &s_t = I_{t-1}(s_{t-1}) \quad \forall t > 0 \\
 &h(I(s_0)) \neq h(s_0)
 \end{aligned} \tag{1}$$

### 3.2 Model architecture

#### 3.2.1 Overall structure

Figure 1 shows the complete FARE model architecture. It is composed of a binary encoder and an RL agent coupled with the Monte Carlo Tree Search procedure. The binary encoder converts the user’s features into a binary representation. The conversion is done by one-hot-encoding the categorical features and discretizing the numerical features into ranges. In the following sections, we will use  $s_t$  to directly indicate the user’s state binary version. Given a state  $s_t$ , the RL agent generates candidate policies,  $\pi_f$  and  $\pi_x$ , for the function and argument generation respectively. MCTS uses these policies as priors for its exploration of the action space and extracts the best next action  $(f, x)_{t+1}^*$ . The action is then applied to the environment. The procedure ends when the STOP action is chosen (i.e., the intervention was successful) or when the maximum intervention length is reached, in which case the result is marked as a failure. During training, the reward is used to improve the MCTS estimates of the policies. Moreover, correct traces (i.e., traces of interventions leading to the desired outcome change) are stored in a replay buffer, and a sample of traces from the buffer is used to refine the RL agent.

**Fig. 3** Agent architecture. Given the user’s state  $s_t$ , it outputs a function policy,  $\pi_f$ , an argument policy  $\pi_x$  and an estimate of the expected reward from the state  $v_t$ . These outputs are used to select the next best action  $(f, x)_{t+1}$



### 3.2.2 RL agent structure

The agent structure is inspired by previous program synthesis works (De Toni et al., 2021; Pierrot et al., 2019). It is composed by 5 components: a state encoder,  $g_{enc}$ , an LSTM controller,  $g_{lstm}$ , a function network  $g_f$ , an argument network  $g_x$  and a value network  $g_v$ . See Fig. 3 for an overview. We use simple feedforward networks to implement  $g_f$ ,  $g_x$  and  $g_v$ .

$$g_{enc}(s_t) = e_t \quad g_{lstm}(e_t, h_{t-1}) = h_t \tag{2}$$

$$g_f(h_t) = \pi_f \quad g_x(h_t) = \pi_x \quad g_v(h_t) = v_t \tag{3}$$

$g_{enc}$  encodes the user’s state in a latent representation which is fed to the controller,  $g_{lstm}$ . The controller,  $g_{lstm}$  learns an implicit representation of the program to generate the interventions. The function and argument networks are then used to extract the corresponding policies,  $\pi_f$  and  $\pi_x$ , by taking as input the hidden state  $h_t$  from  $g_{lstm}$ .  $g_v$  represents the value function  $V$  and it outputs the expected reward from the state  $s_t$ . Here, we omit the state  $s_t$  when defining the policies and the value function output, since  $s_t$  is already embedded into the  $h_t$  representation. In our settings, we try to learn a single program, which we call INTERVENE.

### 3.2.3 Policy

A policy is a distribution over the available actions (i.e., functions and their arguments) such that  $\sum_{i=0}^N \pi(i) = 1$ . Our agent produces two policies:  $\pi_f$  on the function space, and  $\pi_x$  on the argument space. The next action,  $(f, x)_{t+1}$ , is chosen by taking the argmax over the policies:

$$f_{t+1} = \operatorname{argmax}_{f \in \mathcal{F}} \pi_f(f) \quad x_{t+1} = \operatorname{argmax}_{x \in \mathcal{X}_{f_{t+1}}} \pi_x(x|f_{t+1})$$

Each program starts by calling the program INTERVENE, and it ends when the action STOP is called.

### 3.2.4 Reward

Once we have applied the intervention  $I$ , given the black-box classifier  $h$ , the reward,  $r$ , is computed as:

$$r = \lambda^T R \quad \lambda \in (0, 1), R = \begin{cases} 1 & h(I(s)) \neq h(s) \\ 0 & \text{otherwise} \end{cases} \quad (4)$$

where  $\lambda$  is a regularization coefficient and  $T$  is the length of the intervention. The  $\lambda^T$  penalizes longer interventions in favour of shorter ones. Minimizing the intervention length is related to minimizing the *sparsity*, which indicates how many features we have changed to obtain a successful counterfactual (Wachter et al., 2017). *Sparsity* is regarded as an important quality for counterfactual examples and algorithmic recourse (Miller, 2019).

### 3.3 Monte Carlo tree search

Monte Carlo Tree Search (MCTS) is a discrete heuristic search procedure that can successfully solve combinatorial optimization problems with large action spaces (Silver et al., 2018, 2016). MCTS explores the most promising nodes by expanding the search space based on a random sampling of the possible actions. In our setting, each tree node represents the user's state at a time  $t$ , and each arc represents a possible action determining a transition to a new state. MCTS searches for the correct sequence of interventions that minimize the user effort and changes the prediction of the black-box model. We use the agent policies,  $\pi_f$  and  $\pi_x$ , as a prior to explore the program space. Then, the newly found sequence of interventions is used to train the RL agent. To select the next node, we maximize the UCT criterion (Kocsis & Szepesvári, 2006):

$$(f, x)_{t+1} = \operatorname{argmax}_{f \in \mathcal{F}, x \in \mathcal{X}_f} Q(s, (f, x)) + U(s, (f, x)) + L(s, (f, x)) \quad (5)$$

Here  $Q(s, (f, x))$  returns the expected reward by taking action  $(f, x)$ .  $U(s, (f, x))$  is a term that trades-off exploration and exploitation, and it is based on how many times we visited node  $s$  in the tree.  $L(s, (f, x))$  is a scoring term which is defined as follows:

$$L(s, (f, x)) = e^{-(l_{cost}((f,x),s) + l_{count}(f))} \quad (6)$$

where  $l_{cost} = C(a, s) \in \mathbb{R}$  represents the *effort* needed to perform the  $a = (f, x) \in \mathcal{A}$  action, and  $l_{count} \in \mathbb{R}$  penalizes interventions that call multiple times the same function  $f$ . MCTS uses the simulation results to return an improved version of the agent policies  $\pi_f^{mcts}$  and  $\pi_x^{mcts}$ . We can also specify the depth of the search tree as a hyperparameter to balance the computational load requested by the procedure.

From the found intervention, we build an *intervention trace*, which is a sequence of tuples that stores, for each time step  $t$ : the input state, the output state, the reward, the hidden state of the controller and the improved policies. The traces are stored in the replay buffer, to be used to train the RL agent.

### 3.4 Training the agent

The agent has to learn to replicate the interventions provided by MCTS at each step  $t$ . Given the replay buffer, we sample a batch of *intervention traces* and we minimize the cross-entropy  $\mathcal{L}$  between the MCTS policies and the agent policies for each time step  $t$ :

$$\operatorname{argmin}_{\theta} \sum_{\text{batch}} (V - r)^2 - (\pi_f^{\text{mcts}})^T \log(\pi_f) - (\pi_x^{\text{mcts}})^T \log(\pi_x) \quad (7)$$

where  $\theta$  represents the agent's parameters and  $V$  is the value function evaluation computed by the agent.

### 3.5 Generate interventions through RL

When training the agent, we learn a general policy that can be used to provide interventions for many different users. The inference procedure is similar to the one used for training. Given an initial state  $s$ , MCTS explores the tree search space using as “prior” the learnt policies  $\pi_x$  and  $\pi_f$  coming from the agent. The policies  $\pi_x$  and  $\pi_f$  give MCTS a hint of which node to select at each step. Once MCTS finds the minimal cost trace that achieves recourse, we return it to the user. In principle, we can also use only  $\pi_x$  and  $\pi_f$  to obtain a viable intervention (e.g., by deterministically taking the action with highest probability each time). However, keeping the search component (MCTS) with a small exploration budget outperforms the RL agent alone. See Table 2 in Sect. 4 for the comparison between the agent-only model and the agent augmented with MCTS.

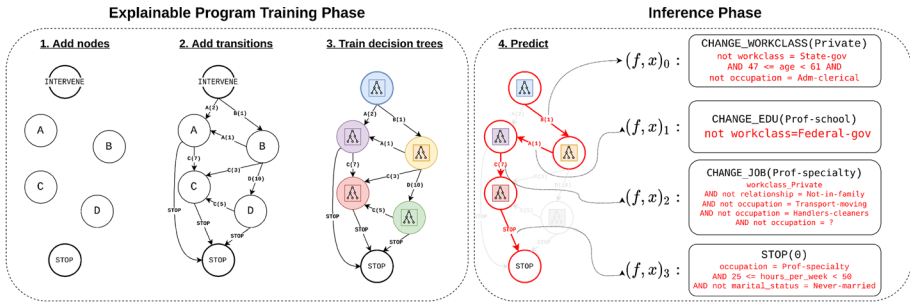
Learning a general policy to provide interventions is a powerful feature. However, the policy is encoded in the latent states of the agent, thus making it impossible for us to understand it. We want to be able to extract from the trained model an explainable version of this policy, which can then be used to *explain* why the model suggested a given intervention. Namely, besides providing to the users a sequence of actions, we want to show also the reason behind each suggested action. The intuition to achieve this is the following: given a set of successful interventions generated by the agent, we can distill a synthetic automaton, or *program*, (E-FARE) which condense the policy in a graph-like structure which we can traverse.

### 3.6 Explainable intervention program

We now show how we can build a deterministic program given the agent. Figure 4 shows the complete procedure and an example of the produced trace. First, we sample  $M$  intervention traces from the trained agent and extract a sequence of  $\{(s_i, (f, x)_i)\}_{i=0}^T$  for each trace. Then, we construct an automaton graph,  $\mathcal{P}$ , in the following way:

1. Given the function library  $\mathcal{F}$ , we create a node for each function  $f$  available. We also add a starting node called INTERVENE and a “sink” node called STOP;
2. We connect each node by unrolling the sampled traces. Starting from INTERVENE, we treat each action  $(f, x)_i$  as a transition. We label the transition with  $(f, x)$  and we connect the current node to the one representing the function  $f$ ;





**Fig. 4** Procedure to generate the explainable program from intervention traces. 1. For all  $f \in \mathcal{F}$ , we add a new node. 2. Given the samples traces, we add the transitions, and we store  $(s_i, (f_i, x_i))$  in each node. 3. We train a decision tree for each node to predict the next action (consistently with the sampled traces). 4. We execute the program on the new instance at prediction time, using the decision trees to decide the next action at each node. We extract a Boolean rule explaining it from the corresponding decision tree for each action. On the right, an example of generated intervention. The actions  $(f, x)$  are black, while the explanations are red (Color figure online)

3. Lastly, for each node  $f$ , we store a collection of outgoing state-action pairs  $(s_i, (f, x)_i)$ . Namely, we store all the states  $s$  and the corresponding outward transitions which were decided by the model while at the node  $f$ ;
4. For each node,  $f \in \mathcal{P}$ , we train a decision tree on the tuples  $(s_i, (f, x)_i)$  stored in the node to predict the transition  $(f, x)_i$  given a user’s state  $s_i$ .

The decision trees are trained only once by using the collection of traces sampled from the trained agent. The agent is frozen at this step, and it is not trained further. At this point, we perform Step 1 to 3 of Fig. 4. The pseudocode of the entire procedure is available in the ["Appendix A"](#).

### 3.7 Generate explainable interventions

The intervention generation is done by traversing the graph  $\mathcal{P}$ , starting from the node INTERVENE, until we reach the STOP node or we reach the maximum intervention length. In the last case, the program is marked as a failure. Given the node  $f \in \mathcal{P}$  and given the state  $s_t$ , we use the decision tree of that node to predict the next transition  $(f', x')$ . Moreover, we can extract from the decision tree interpretable rules which tell us why the next action was chosen. A rule is a boolean proposition on the user’s features such as  $(income > 5000 \wedge education = bachelor)$ . Then, we follow  $(f', x')$ , which is an arc going from  $f$  to the next node  $f'$ , and we apply the action to  $s_t$  to get  $s_{t+1}$ . Again, the program is “fixed” at inference time, and it is not trained further. See Step 4 of Fig. 4 for an example of the inference procedure and of the produced explainable trace.

**Table 1** Description of the datasets

Dataset	$ D $	$h(s) = 1$	$h(s) = 0$	$ s $	$ B(s) $	$ F $
<i>german</i>	1002	301	701	10	44	7
<i>adult</i>	48,845	11,691	37,154	15	125	6
<i>syn</i>	10,004	5002	5002	10	40	6
<i>syn_long</i>	10,004	5002	5002	14	64	10

$|D|$  is the size of the dataset.  $|s|$  the number of features for an instance.  $|B(s)|$  shows how many binary features the agent sees after the conversion with the binary converter.  $|F|$  is the size of the agent program library.  $h(s)$  indicates the number of favourable (1) and unfavourable (0) samples

## 4 Experiments


Our experimental evaluation aims at answering the following research questions: (1) Does our method provide better performances than the competitors in terms of the *validity* of the algorithmic recourse? (2) Does our approach allow us to complement interventions with action-by-action explanations in most cases? (3) Does our method minimize the interaction with the black-box classifier to provide interventions?

The code and the dataset of the experiments are available on Github to ensure reproducibility.<sup>1</sup> The software exploit parallelization through `mpi4python` (Dalcin & Fang, 2021) to improve inference and training time. We compared the performance of our algorithm with CSCF (Naumann & Ntoutsis, 2021), to the best of our knowledge the only existing model-agnostic approach that can generate consequence-aware interventions following a causal graph. However, note that earlier solutions still perform user-specific optimization, so that our results in terms of generality, interpretability and cost (number of queries to the black-box classifier and computational cost) carry over to these alternatives. For the sake of a fair comparison, we built our own parallelized version of the CSCF model based on the original code. We developed the project to make it easily extendable and reusable by the research community. The experiments were performed using a Linux distribution on an Intel(R) Xeon(R) CPU E5-2660 2.20GHz with 8 cores and 100 GB of RAM (only 4 cores were used).

### 4.1 Dataset and black-box classifiers

Table 1 shows a brief description of the datasets. They all represent binary (favourable/unfavourable) classification problems. The two real world datasets, *German Credit* (*german*) and *Adult Score* (*adult*) (Dua & Graff, 2017), are taken from the relevant literature. Given that in these datasets a couple of actions is usually sufficient to overturn the outcome of the black-box classifier, we also developed two synthetic datasets, *syn* and *syn\_long*, where longer interventions are required, so as to evaluate the models in more challenging scenarios. The datasets are made of both categorical and numerical features (e.g., monthly income, job type, etc.). Each dataset was randomly split into 80% train and 20% test. For each dataset, we manually define a causal graph,  $\mathcal{G}$ , by looking at the features available.

<sup>1</sup> <https://github.com/unitn-sml/syn-interventions-algorithmic-recourse>.



**Fig. 5** Experimental results. (Left) *validity* (fraction of successful interventions); (Middle) Average length of a successful intervention; (Right) Average cost of a successful intervention. Results are averaged over 100 test examples

For the synthetic datasets, we sampled instances directly from the causal graph. See Fig. 10 in the Appendix for an example of these graphs. The black-box classifier for *german* and *adult* was obtained by training a 5-layers MLP with ReLU activations. The trained classifiers are reasonably accurate ( $\sim 0.9$  test-set accuracy for *german*,  $\sim 0.8$  for *adult*). The synthetic datasets (*syn* and *syn\_long*) do not require any training since we directly use our manually defined decision function.

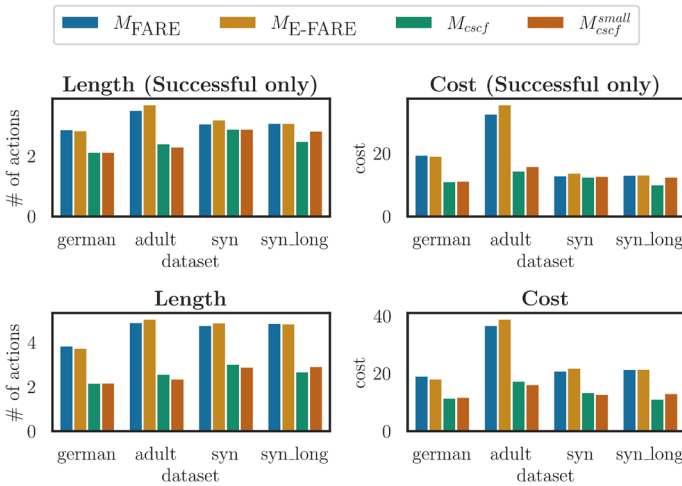
## 4.2 Models

We evaluate four different models: FARE, the agent coupled with MCTS ( $M_{FARE}$ ), E-FARE, the explainable deterministic program distilled from the agent ( $M_{E-FARE}$ ), and two versions of CSCF, one ( $M_{cscf}$ ) with a large budget of generation,  $n$ , and population size,  $p$ , ( $n = 50, p = 200$ ) and one ( $M_{cscf}^{small}$ ) with a smaller budget ( $n = 25, p = 100$ ). For  $M_{FARE}$ , we set the MCTS exploration depth to 7 for all the experiments.

## 4.3 Evaluation

The left plot in Fig. 5 shows the average *validity* of the different models, namely the fraction of instances for which a model manages to generate a successful intervention (Wachter et al., 2017). We can see how  $M_{FARE}$  outperforms or is on-par with the  $M_{cscf}$  and  $M_{cscf}^{small}$  models on both the real-world and synthetic datasets. The performance difference is more evident in the synthetic datasets because the evolutionary algorithm struggles to generate interventions that require more than a couple of actions. The *validity* loss incurred in distilling  $M_{FARE}$  into a program ( $M_{E-FARE}$ ) is rather limited. This implies that we are able to provide interventions with explanations for 94% (*german*), 66% (*adult*), 99% (*syn*) and 87% (*syn\_long*) of the test users.<sup>2</sup> Moreover,  $M_{E-FARE}$  generates similar interventions to  $M_{FARE}$ . The sequence similarity between their respective interventions for the same user

<sup>2</sup> Note that the *validity* loss observed on *adult* is due to the limited sampling budget we allocated for  $M_{E-FARE}$  (250 traces for all datasets). Adapting this budget to the feature space size (considerably larger for *adult*) can help boost the performance, at the cost of generating longer explanations.



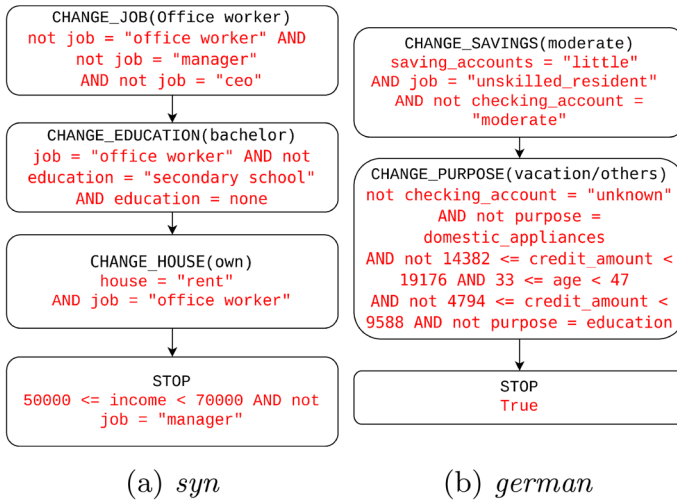
**Fig. 6** Evaluation considering only the instances for which all the models provide a successful intervention. If we restrict the comparison to the subset of instances for which all models manage to generate a successful intervention, the difference in costs between methods shrinks substantially (top left vs bottom left). The same behaviour applies to the intervention length (top right vs bottom right)

are 0.89 (*german*), 0.72 (*adult*), 0.80 (*syn*) and 0.71 (*syn\_long*), where 1.0 indicates identical interventions.

The main reason for the *validity* gains of our model is the ability to generate long interventions, something evolutionary-based algorithms struggle with. This effect can be clearly seen from the middle plot of Fig. 5. Both  $M_{cscf}$  and  $M_{cscf}^{small}$  rarely generate interventions with more than two actions, while our approach can easily generate interventions with up to five actions. A drawback of this ability is that intervention costs are, on average, higher (right plot of Fig. 5). On the one hand, this is due to the fact that our model is capable of finding interventions for more complex instances, while  $M_{cscf}$  and  $M_{cscf}^{small}$  fail. Indeed, if we compute lengths and costs on the subset of instances for which all models find a successful intervention, the difference between the approaches is less pronounced. See Fig. 6 for the evaluation. On the other hand, there is a clear trade-off between solving a new optimization problem from scratch for each new user, and learning a general model that, once trained, can generate interventions for new users in real-time and without accessing the black-box classifier.

We also conducted a quantitative analysis of the quality of the explanations generated using  $M_{E-FARE}$ . We measured the average number of boolean clauses in the rule of a given suggested action. Our explanations need to be concise, thus involving a limited number of features, to be easily understandable. The literature defines seven as the maximum acceptable number of concepts in an explanation (Miller, 2019, 1956). We have an average of 3 for the *syn*, *syn\_long* and *german* datasets, while we have an average of 6.5 clauses for the *adult* dataset. Indeed, the  $M_{E-FARE}$  model can generate compact explanations as boolean predicates. The *adult* dataset requires a more complex recourse policy. Therefore the decision rules of the automaton are more complex, thus involving longer boolean predicates. See Fig. 7 for examples of interventions coupled with rule-based explanations.

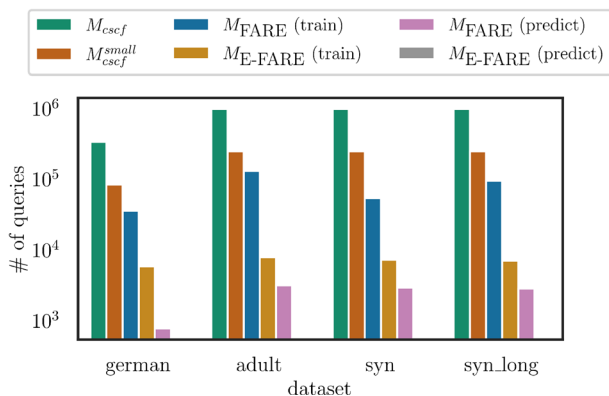
Figure 8 reports the average number of queries to the black-box classifier. Our approach requires far fewer queries than  $M_{cscf}$  (note that the plot is in logscale), and



**Fig. 7** Example of Interventions with rule-based explanations. We show here two additional examples of successful interventions (*syn* and *german* datasets) combined with boolean predicates explaining why we suggested the given action. The black text indicates the action ( $f, x$ ), while the red text indicates the decision rule (Color figure online)

even substantially less than  $M_{cscf}^{small}$  (that is anyhow not competitive in terms of *validity*). Furthermore, most queries are made for training the agent ( $M_{FARE}(train)$ ), which is only done once for all users. Once the model is trained, generating interventions for a single user requires around two orders of magnitude fewer queries than the competitors. Note that MCTS is crucial to allow the RL agent to learn a successful policy with a low budget of queries. Indeed, training an RL agent without the support of MCTS fails to converge in the given budget (between 50 and 100 iterations), leading to a completely useless policy. By efficiently searching the space of interventions, MCTS manages to quickly correct inaccurate initial policies, allowing the agent to learn high quality policies with a limited query budget. MCTS is also critical during inference, since it increases the *validity* of the results. Given a trained agent, the *validity* drops if we perform inference without the MCTS components. See Table 2 for the evaluation.

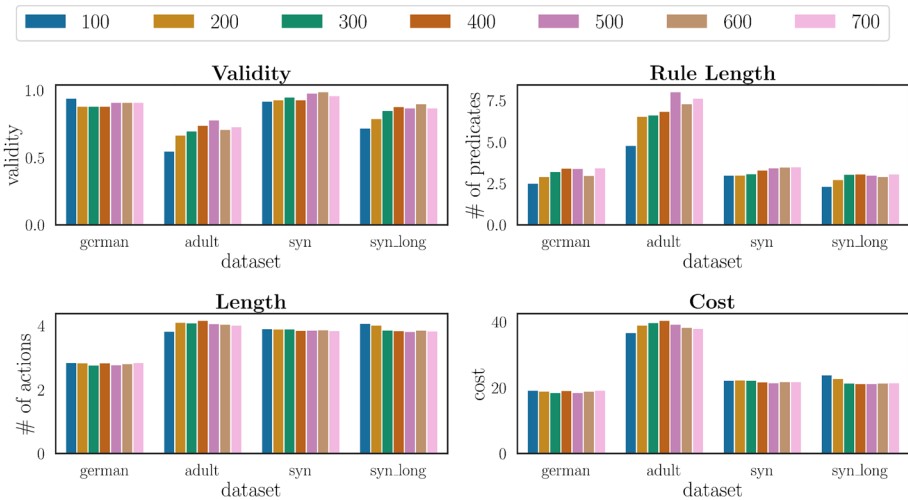
**Fig. 8** Number of queries. Total number of queries to the black-box classifier made by the models.  $M_{E-FARE}(predict)$  is not visible, as the automaton does not query the black-box classifier to generate interventions. Note that the number of queries is in logscale



**Table 2** Ablation study

Dataset	$M_{\text{FARE}}$	$M_{\text{agent}}$
<i>german</i>	<b>1.00</b>	0.00
<i>adult</i>	<b>0.93</b>	0.00
<i>syn</i>	<b>0.98</b>	0.59
<i>syn_long</i>	<b>0.92</b>	0.00

In order to evaluate the contribution of MCTS in finding successful interventions, we also evaluated a trained FARE model which only uses the RL agent,  $M_{\text{agent}}$ . The agent predicts the next action using its own policy without leveraging MCTS to refine the choice. Results indicate that RL alone is incapable of finding successful interventions and the *validity* drops. Bold values indicate the best results



**Fig. 9** Validity of  $M_{\text{E-FARE}}$  when varying the training budget. We show the effect on increasing the sampling budget (from 100 to 700 traces) when training the  $M_{\text{E-FARE}}$  model

When turning to the program, building the automaton ( $M_{\text{E-FARE}}(\text{train})$ ) requires a negligible number of queries to extract the intervention traces used as supervision.

Using the automaton to generate interventions *does not* require to query the black-box classifier. This characteristic can substantially increase the usability of the system, as  $M_{\text{E-FARE}}$  can be employed directly by the user even if they have no access to the classifier. Computationally speaking, the advantage of a two-step phase is also quite dramatic.  $M_{\text{cscf}}$  takes an average of  $\sim 693$  s for each user to provide a solution (the same order of magnitude of training a model for all users with  $M_{\text{FARE}}$ ), while  $M_{\text{FARE}}$  inference time is under 1s, allowing real-time interaction with the user.

Additionally, Fig. 9 shows how it is possible to improve the performances of  $M_{\text{E-FARE}}$  by just sampling more traces from the trained agent ( $M_{\text{FARE}}$ ). We can see how the *validity* increases in the *adult*, *syn* and *syn\_long* datasets. We also notice that using a larger budget to train  $M_{\text{E-FARE}}$  produces longer explainable rules by keeping the length and cost of the generated interventions almost constant. The total number of queries to the black-box classifier will also slightly increase.

Overall, our experimental evaluation allows us to affirmatively answer the research questions stated above.

## 5 Conclusion

This work improves the state-of-the-art on algorithmic recourse by providing a method, FARE (eFficient counterfActual REcourse), that can generate effective and interpretable counterfactual interventions in real-time. Our experimental evaluation confirms the advantages of our solution with respect to alternative consequence-aware approaches in terms of *validity*, interpretability and number of queries to the black-box classifier. Our work unlocks many new research directions, which could be explored to solve some of its limitations. First, following previous work on causal-aware intervention generation, we use manually-crafted causal graphs and action costs. Learning them from the available data directly, minimizing the human intervention, would allow applying the approach in settings where this information is not available or unreliable. Second, we showed how our method learns a general program by optimizing over multiple users. It would be interesting to investigate additional RL methods to optimize the interventions globally and locally to provide more personalized sequences to the users. Such methods could be coupled with interactive approaches eliciting preferences and constraints directly from the user, thus maximizing the chance to generate the most appropriate intervention for a given user.

## 6 Ethical Impact

The research field of algorithmic recourse aims at improving fairness, by providing unfairly treated users with tools to overturn unfavourable outcomes. By providing real-time, explainable interventions, our work makes a step further in making these tools widely accessible. As for other approaches providing counterfactual interventions, our model could in principle be adapted by malicious users to “hack” a fair system. Research on adversarial training can help in mitigating this risk.

## Appendix A: Program distillation pseudocode

We present here the pseudocode of two algorithms. Algorithm 1 shows how to distill the synthetic program from the agent and it refers to Step 3 of Fig. 4. Algorithm 1 shows how the distilled program is applied at inference time to a new user and it refers to Step 4 of Fig. 4.

## Appendix B: Domain specific languages (DSL) and causal graphs

We now show the Domain Specific Languages (DSL) used for the *german* (Table 3), *synthetic* (Table 4) and *adult* (Table 5) experiments. For each setting, we show the functions available, the argument type they accept and an exhaustive list of the potential arguments. Each program operates on a single feature. The name of the program suggests the feature it operates on (e.g., *CHANGE\_JOB* operate on the feature *job*). The programs which accept

**Table 3** Example of the DSL for the *german* experiment

Program	Argument Type	Argument
CHANGE_SAVINGS	Categorical	unknown, little, moderate, rich, quite_rich
CHANGE_JOB	Categorical	unskilled_non_resident, unskilled_resident, skilled, highly_skilled
CHANGE_CREDIT	Numerical	100, 1000, 2000, 5000
CHANGE_HOUSING	Categorical	free, rent, own
CHANGE_DURATION	Numerical	10, 20, 30
CHANGE_PURPOSE	Categorical	business, car, domestic_appliances, education, furniture/equipment, radio/TV, repairs, vacation/others
STOP	–	–

**Table 4** Example of the DSL for the *synthetic* experiment

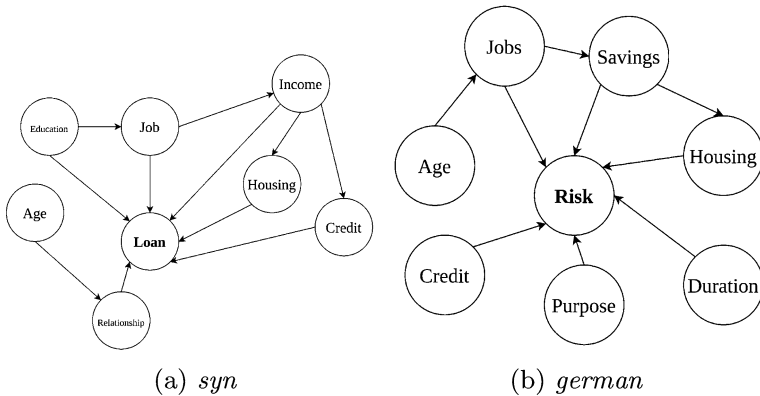
Program	Argument Type	Argument
CHANGE_EDUCATION	Categorical	none,secondary school diploma, bachelor, master, phd
CHANGE_JOB	Categorical	unemployed, worker, office worker, manager, ceo
CHANGE_INCOME	Numerical	5000, 10000, 20000, 30000, 40000, 50000
CHANGE_HOUSE	Categorical	none, rent, own
CHANGE_RELATION	Categorical	single, married, divorced, widow/er
STOP	–	–

numerical arguments simply add their argument to the current value of the target feature. The program *STOP* does not accept any argument and it signals only the end of the intervention without changing the features. The DSLs for the *synthetic\_long* experiment is similarly defined and is omitted for brevity (Fig. 10).



**Table 5** Example of the DSL for the *adult* experiment

Program	Argument Type	Argument
CHANGE_WORKCLASS	Categorical	Never-worked, Without-pay, Self-emp-not-inc, Self-emp-inc, Private, Local-gov, State-gov, Federal-gov, ?
CHANGE_EDUCATION	Categorical	Preschool, 1st-4th, 5th-6th, 7th-8th, 9th, 10th, 11th, 12th, HS-grad, Some-college, Bachelors, Masters, Doctorate, Assoc-acdm, Assoc-voc, Prof-school
CHANGE_OCCUPATION	Numerical	Tech-support, Craft-repair, Other-service, Sales, Exec-managerial, Prof-specialty, Handlers-cleaners, Machine-op-inspect, Adm-clerical, Farming-fishing, Transport-moving, Priv-house-serv, Protective-serv, Armed-Forces, ?
CHANGE_RELATIONSHIP	Categorical	Wife, Own-child, Husband, Not-in-family, Other-relative, Unmarried
CHANGE_HOURS	Numerical	Range from 0 to 25
STOP	—	—



**Fig. 10** Causal Graphs. Depiction of the causal graphs used in the experiments for the *syn* and *german* dataset. The bold nodes indicate the variables we want to predict. Here, the graphs encode the assumption that we know the factors influencing the target features (**Risk** and **Loan**). However, in practice, we cannot know which features the decision function of the black-box model is using for inference

---

**Algorithm 1** Generate the explainable program. Given the automaton reconstructed from the traces  $\hat{\mathcal{P}}$ , we generate the explainable program  $\mathcal{P}$ . For each node, we train the decision tree classifier on the traces stored  $(s, (f, x))_i$  to predict  $(f, x)$  from  $s$ .

---

**Input:**  $\hat{\mathcal{P}}$ , automaton extracted from the traces

**Output:**  $\mathcal{P}$ , explainable program

```

1: Let  $t = 0$ .
2: Let  $\mathcal{P} = \{\}$ 
3: for all  $node \in \hat{\mathcal{P}}$  do
4:    $program = node.program$ 
5:   if  $program \neq STOP$  then
6:     if  $node.arcs.unique() > 1$  then
7:       Train a decision tree,  $\mathcal{D}$ , on  $node.arcs$ 
8:        $node.classifier = \mathcal{D}$ 
9:     else
10:       $node.classifier = \{True \rightarrow node.arcs[0]\}$ 
11:    end if
12:     $\mathcal{P}.append(node)$ 
13:  end if
14: end for
15: return  $\mathcal{P}$ 

```

---

---

**Algorithm 2 Predict with the program.** Given the program  $\mathcal{P}$ , we can infer the complete explainable intervention for a user  $s_0$  by traversing  $\mathcal{P}$ . This can be done recursively. The recursion ends when we either predict the STOP action or we reach the maximum intervention length,  $\alpha \in \mathbb{N}^+$ .

---

**Input:**  $\mathcal{P}, s_0, (f, x)_0, \mathcal{I}, \mathcal{R}$

**Output:**  $s$ , the final state,  $\mathcal{I}$ , the intervention actions,  $\mathcal{R}$ , rules for each action

```

1: if  $(f, x)_t == (STOP, 0) \vee t > \alpha$  then
2:   return  $s_{t+1}, \mathcal{I}, \mathcal{R}$ 
3: else
4:    $(f, x)_{t+1}, rule_{t+1} = \mathcal{P}.get(f).predict(s_t)$ 
5:    $s_{t+1} = (f, x)_t.apply(s_t)$ 
6:
7:    $\mathcal{R}.append(rule_{t+1})$ 
8:    $\mathcal{I}.append((f, x)_{t+1})$ 
9:   recursive call  $(\mathcal{P}, s_{t+1}(f, x)_{t+1}, \mathcal{I}, \mathcal{R})$ 
10: end if

```

---

**Author contributions** GDT designed the method, conducted the data collection process, built the experimental infrastructure and performed the relevant experiments. BL and AP contributed to the design of the method, provided supervision and resources. All authors contributed to the writing of the manuscript.

**Funding** This research was partially supported by TAILOR, a project funded by EU Horizon 2020 research and innovation programme under GA No 952215. The work of Giovanni De Toni was partially supported by the project AI@Trento (FBK-Unitn).

**Availability of data and materials** The datasets used in the experimental evaluation are freely available at <https://github.com/unitn-sml/syn-interventions-algorithmic-recourse>.

**Code availability** The code is freely available at <https://github.com/unitn-sml/syn-interventions-algorithmic-recourse>.

## Declarations

**Conflict of interest** The authors have no competing interests to declare that are relevant to the content of this article.

**Ethics approval** Not applicable.

**Consent to participate** Not applicable.

**Consent for publication** Not applicable.

**Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not

permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

## References

- Apley, D. W., & 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.
- Dalcin, L., & Fang, Y. (2021). mpi4py: Status update after 12 years of development. *Computing in Science Engineering*, 23(4), 47–54. <https://doi.org/10.1109/MCSE.2021.3083216>.
- Miller, G. A. (1956). The magical number seven, plus or minus two: Some limits on our capacity for processing information. *Psychological Review*, 63(2), 81.
- Miller, T. (2019). Explanation in artificial intelligence: Insights from the social sciences. *Artificial Intelligence*, 267, 1–38.
- Pearl, J. (2009). *Causality*. Cambridge University Press.
- Pierrot, T., Ligner, G., Reed, S. E., Sigaud, O., Perrin, N., Laterre, A., et al. (2019). Learning compositional neural programs with recursive tree search and planning. *NeurIPS*, 32, 14673–14683.
- Silver, D., Huang, A., Maddison, C. J., Guez, A., Sifre, L., van den Driessche, G., et al. (2016). Mastering the game of go with deep neural networks and tree search. *Nature*, 529, 484–503.
- Silver, D., Hubert, T., Schrittwieser, J., Antonoglou, I., Lai, M., Guez, A., et al. (2018). A general reinforcement learning algorithm that masters chess, shogi, and go through self-play. *Science*, 362(6419), 1140–1144. <https://doi.org/10.1126/science.aar6404>.
- Stepin, I., Alonso, J. M., & Pereira-Fariña, A. C. M. (2021). A survey of contrastive and counterfactual explanation generation methods for explainable artificial intelligence. *IEEE Access*, 9, 11974–12001.
- Voigt, P., & Bussche, A. (2017). *The EU general data protection regulation (GDPR): A practical guide* (1st ed.). Springer.
- Wachter, S., Mittelstadt, B., & Russell, C. (2017). Counterfactual explanations without opening the black box: Automated decisions and the GDPR. *Harvard Journal of Law and Technology*, 31, 841.
- Balog, M., Gaunt, A. L., Brockschmidt, M., Nowozin, S., & Tarlow, D. (2017). DeepCoder: Learning to write programs. In *ICLR*. <https://openreview.net/pdf?id=rkE3y85ee>
- Barocas, S., Selbst, A., & Raghavan, M. (2020). The hidden assumptions behind counterfactual explanations and principal reasons. In *FAT\**.
- Bunel, R., Hausknecht, M., Devlin, J., Singh, R., & Kohli, P. (2018). Leveraging grammar and reinforcement learning for neural program synthesis. In *ICLR*. <https://openreview.net/forum?id=H1Xw62kRZ>
- Coulom, R. (2006). Efficient selectivity and backup operators in Monte-Carlo tree search. In *Proceedings computers and games 2006*. Springer.
- Cui, Z., Chen, W., He, Y., & Chen, Y. (2015). Optimal action extraction for random forests and boosted trees. In *Proceedings of the 21th ACM SIGKDD international conference on knowledge discovery and data mining* (pp. 179–188).
- Dandl, S., Molnar, C., Binder, M., & Bischl, B. (2020). Multi-objective counterfactual explanations. In *PPSN* (pp. 448–469). Springer.
- De Toni, G., Erculiani, L., & Passerini, A. (2021). Learning compositional programs with arguments and sampling. In *AIPLANS*.
- Dua, D., & Graff, C. (2017). UCI machine learning repository. <http://archive.ics.uci.edu/ml>
- Greenwell, B. M., Boehmke, B. C., & McCarthy, A. J. (2018). A simple and effective model-based variable importance measure. arXiv preprint [arXiv:1805.04755](https://arxiv.org/abs/1805.04755)
- Guidotti, R., Monreale, A., Ruggieri, S., Pedreschi, D., Turini, F., & Giannotti, F. (2018). Local rule-based explanations of black box decision systems. CoRR [arxiv:1805.10820](https://arxiv.org/abs/1805.10820)
- Kanamori, K., Takagi, T., Kobayashi, K., & Arimura, H. (2020). Dace: Distribution-aware counterfactual explanation by mixed-integer linear optimization. In *IJCAI* (pp. 2855–2862).
- Karimi, A., Barthe, G., Balle, B., & Valera, I. (2020). Model-agnostic counterfactual explanations for consequential decisions. In *AISTATS* (pp. 895–905). PMLR.
- Karimi, A., Barthe, G., Schölkopf, B., & Valera, I. (2020). A survey of algorithmic recourse: Definitions, formulations, solutions, and prospects. arXiv preprint [arXiv:2010.04050](https://arxiv.org/abs/2010.04050)
- Karimi, A., Schölkopf, B., & Valera, I. (2021). Algorithmic recourse: from counterfactual explanations to interventions. In *FaccT* (pp. 353–362).

- Karimi, A., von Kügelgen, J., Schölkopf, B., & Valera, I. (2020). Algorithmic recourse under imperfect causal knowledge: A probabilistic approach. In *NeurIPS*. <https://proceedings.neurips.cc/paper/2020/file/02a3c7fb3f489288ae6942498498db20-Paper.pdf>
- Kocsis, L., & Szepesvári, C. (2006). Bandit based Monte-Carlo planning. In *ECML* (pp. 282–293). Springer, Berlin, Heidelberg. [https://doi.org/10.1007/11871842\\_29](https://doi.org/10.1007/11871842_29)
- Mothilal, R. K., Sharma, A., & Tan, C. (2020). Explaining machine learning classifiers through diverse counterfactual explanations. In *FAT\** (pp. 607–617).
- Naumann, P., & Ntoutsis, E. (2021). Consequence-aware sequential counterfactual generation. In *ECMLP-KDD*. [https://doi.org/10.1007/978-3-030-86520-7\\_42](https://doi.org/10.1007/978-3-030-86520-7_42)
- Ramakrishnan, G., Lee, Y. C., & Albarghouthi, A. (2020). Synthesizing action sequences for modifying model decisions. In *AAAI* (Vol. 34, pp. 5462–5469).
- Spirtes, P., & Zhang, K. (2016). Causal discovery and inference: concepts and recent methodological advances. In *Applied Informatics* (Vol. 3, pp. 1–28). SpringerOpen.
- Tian, J., & Pearl, J. (2001). Causal discovery from changes. In *Proceedings of the seventeenth conference on uncertainty in artificial intelligence* (pp. 512–521).
- Tolomei, G., Silvestri, F., Haines, A., & Lalmas, M. (2017). Interpretable predictions of tree-based ensembles via actionable feature tweaking. In *Proceedings of the 23rd ACM SIGKDD international conference on knowledge discovery and data mining* (pp. 465–474).
- Tsirtsis, S., & Rodriguez, M. (2020). Decisions, counterfactual explanations and strategic behavior. In *NeurIPS*. <https://proceedings.neurips.cc/paper/2020/hash/c2ba1bc54b239208cb37b901c0d3b363-Abstract.html>
- Ustun, B., Spangher, A., & Liu, Y. (2019). Actionable recourse in linear classification. In *FAT\** (pp. 10–19).
- Yonadav, S., & Moses, W. S. (2019). Extracting incentives from black-box decisions. CoRR [arxiv:1910.05664](https://arxiv.org/abs/1910.05664)

**Publisher's Note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.