



# Preference Elicitation in Interactive and User-centered Algorithmic Recourse: An Initial Exploration

Seyedehdelaram Esfahani  
delaramesfahani73@gmail.com  
University of Trento  
Italy

Giovanni De Toni  
giovanni.detoni@unitn.it  
University of Trento & Fondazione  
Bruno Kessler  
Italy

Bruno Lepri  
lepri@fbk.eu  
Fondazione Bruno Kessler  
Italy

Andrea Passerini  
andrea.passerini@unitn.it  
University of Trento  
Italy

Katya Tentori  
Massimo Zancanaro  
katya.tentori@unitn.it  
massimo.zancanaro@unitn.it  
University of Trento & Fondazione  
Bruno Kessler  
Italy

## ABSTRACT

Algorithmic Recourse aims to provide actionable explanations, or recourse plans, to overturn potentially unfavourable decisions taken by automated machine learning models. In this paper, we propose an interaction paradigm based on a *guided* interaction pattern aimed at both eliciting the users' preferences and heading them toward effective recourse interventions. In a fictional task of money lending, we compare this approach with an *exploratory* interaction pattern based on a combination of alternative plans and the possibility of freely changing the configurations by the users themselves. Our results suggest that users may recognize that the *guided* interaction paradigm improves efficiency. However, they also feel less freedom to experiment with "what-if" scenarios. Nevertheless, the time spent on the purely *exploratory* interface tends to be perceived as a lack of efficiency, which reduces attractiveness, perspicuity, and dependability. Conversely, for the *guided* interface, more time on the interface seems to increase its attractiveness, perspicuity, and dependability while not impacting the perceived efficiency. That might suggest that this type of interfaces should combine these two approaches by trying to support *exploratory* behavior while gently pushing toward a *guided* effective solution.

## CCS CONCEPTS

• **Human-centered computing** → **HCI design and evaluation methods**; • **Computing methodologies** → **Machine learning**.

## KEYWORDS

Human-centred AI, Algorithmic Recourse, Counterfactual Examples

### ACM Reference Format:

Seyedehdelaram Esfahani, Giovanni De Toni, Bruno Lepri, Andrea Passerini, Katya Tentori, and Massimo Zancanaro. 2024. Preference Elicitation in Interactive and User-centered Algorithmic Recourse: An Initial Exploration. In *Proceedings of the 32nd ACM Conference on User Modeling, Adaptation and Personalization (UMAP '24)*, July 01–04, 2024, Cagliari, Italy. ACM, New York, NY, USA, 6 pages. <https://doi.org/10.1145/3627043.3659556>

## 1 INTRODUCTION

Alice is applying for a loan at a respected bank in her hometown that uses an automated machine learning system to decide whether to accept or reject loan applicants. Unfortunately, Alice gets denied the loan and the bank cannot provide her with any explanation about their decision, since their system is a "black-box": it performs well, but it is too complicated for humans to understand.

Alice's story is fictional, but there are already examples of machine learning systems applied to decision-making tasks with a high impact on human lives, such as credit scoring [2] and recidivism prediction [32]. Indeed, more than just an explanation about the reasons for the rejection of her request, Alice would benefit from practical suggestions on how to reverse the algorithm decision. *Algorithmic recourse* (AR) [13, 31] is a particular case of explanations based on counterfactual analysis [12] that can provide a sequence of actions capable of overturning an unfavorable machine-driven decision [27]. For example, a recourse suggestion would provide Alice with the actions she needs to take to obtain a loan (e.g., "increase your monthly salary by 100\$"). In this respect, algorithmic recourse can better respond to the need for transparency and fairness required for AI-based systems by several national authorities in a context where AI-based tools can significantly impact people's lives [31]. This aspect is gaining increasing importance as one effective way to implement the rights sought by the European GDPR [30].

Although there is a growing body of research on this topic [13, 28], the assessment of the state-of-the-art recourse algorithms

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

UMAP '24, July 01–04, 2024, Cagliari, Italy

© 2024 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 979-8-4007-0433-8/24/07

<https://doi.org/10.1145/3627043.3659556>

is usually based on quantitative benchmarking with metrics such as validity, minimality or similarity [3]. While helpful in appreciating the potential of algorithmic recourse strategies, these approaches fall short of assessing the effectiveness of recourse strategies for real users. While an algorithm must be efficient and effective (as well as free for bias[8]), it is also of paramount importance that the interaction paradigm sustains and fosters empowerment by providing a creative space of opportunities to explore while limiting confusion and false expectations. In this direction, the recent work by Wang and colleagues [33] proposes an algorithm based on integer linear programming to generate counterfactual explanations, together with a graphical user interface called GAM Coach, to enable end users to develop recourse plans iteratively. The authors refer to their approach as an *exploratory interface* by citing the seminal work by Shneiderman on Human-centered AI Systems [24, 25] since it allows users to freely manipulate the values of the features involved in the decision process. An evaluation study of GAM Coach provided helpful insights on designing user interfaces to support users in leveraging these algorithms effectively. Nevertheless, the authors noted how the exploration might become complicated and require several interactions to learn how it works. As a design lesson, they suggest that “developers can use the log data [...] to train a new model to predict users’ preference configurations” [33]. Another interesting result of that study is that, although participants appreciate being able to control recourse generation with what-if questions, “interactivity and transparency could occasionally confuse users with counterintuitive recourse plans”. [33]. Indeed, a less explored but critical desideratum is the degree of *personalization* of the recourse plans to match the users’ needs. Tailoring recourse to users requires an *interactive* strategy to include humans in the generation process since some users’ preferences cannot be inferred by their features alone [26].

Our work investigates a different interaction paradigm with respect to the *purely exploratory* approach of GAM Coach, while still aimed at implementing the *proactive principle* proposed for Human-centered AI Systems: “to probe the algorithm boundaries with different inputs” [24]. Specifically, we use an algorithmic recourse method enriched with preference elicitation capabilities which realizes the *personalization* of a recourse plan. We propose an approach based on a *guided-interaction* pattern aimed at directly eliciting the user’s preferences needed by the algorithm to head the user toward a satisfying recourse intervention quickly. We believe that limiting the exploratory opportunities within an optimized *step-by-step* process may help relieve the user of the possible overwhelming amount of choices in a typical recourse negotiation while still preserving the gist of the prospective strategy by providing users with “a better understanding of each step in the process so that they can prevent mistakes” [25].

In this regard, our primary research question concerns the user experience of a novel interaction strategy that we refer to as *guided interaction*. This strategy maintains an exploratory interaction style while establishing boundaries via an optimized preference elicitation process.

*RQ:* To what extent do users recognize and appreciate a *guided-interaction* strategy as measured by the

pragmatic and hedonic dimensions of their overall experience?

To provide a first answer to our research question, we present the design of a graphical user interface based on the *guided-interaction* pattern as the foreground of a novel preference elicitation algorithm for recourse. Then, we compare it with a pilot study in a typical scenario for recourse problems (a fictional task of money lending) to a *purely exploratory* interface inspired by the GAM Coach one [33].

## 2 RELATED WORK

Algorithmic Recourse is concerned with providing actionable suggestions, or *interventions*, that show how to overturn an unfavorable prediction of a machine learning model via counterfactual scenarios [6, 31]. Formally, let a user sampled from a distribution  $P(X)$  be  $\mathbf{x} = \{x_1, \dots, x_d\} \in \mathcal{X}$ , where  $x_i$  is a single feature (e.g., age, education level, etc.). Each user  $x$  is assigned a binary outcome  $y \in \{0, 1\}$  following  $P(Y | X)$ . Given a dataset  $\{(x, y)_i\}_{i=1}^N$  sampled from  $P(X, Y)$ , we can learn a classifier  $h : \mathcal{X} \rightarrow \{0, 1\}$  which predicts the class  $y$  given a user  $x$ . Following our previous example,  $y = 1$  could mean “ $x$  will repay the loan”. We also define the (possibly unknown) user preferences as a vector  $\mathbf{w} \in \mathcal{W}$  sampled from a distribution  $P(W)$ . An intervention is a sequence of actions  $I = \{a^{(1)}, \dots, a^{(N)}\}$  where each action  $a^{(i)} \in \mathcal{A}$  suggests a modification to a single *actionable* feature e.g., “reduce your monthly spending by 100\$”. Given the classifier  $h : \mathcal{X} \rightarrow \{0, 1\}$ , in algorithmic recourse we aim at finding the best sequence of actions  $I$  optimizing the following objective:

$$I^* = \underset{I}{\operatorname{argmin}} c(I, \mathbf{x}, \mathbf{w}) \quad \text{s.t.} \quad h(I(\mathbf{x})) \neq h(\mathbf{x}) \quad (1)$$

where  $c : \mathcal{I} \times \mathcal{X} \times \mathcal{W} \rightarrow \mathbb{R}^+$  is a cost function. In practice, the idea is to find the *cheapest* intervention following the user preferences that, if applied to the initial profile, will induce a different outcome by the classifier. Optimizing eq. 1 means solving both a *NP-hard* combinatorial problem together with estimating the user preferences. There is already a plethora of methods for optimizing eq. 1 [13, 28], which however discard user preferences or make some unrealistic assumptions. Many algorithmic recourse approaches assume preferences are fully specified beforehand [10, 14, 19, 31] or they assume users can define them via numerical constraints [18] or by quantifying the cost of each action [16]. However, it is well known that it is hard for users to specify preferences quantitatively [15].

For what concerns user interaction and assessment of user experience, recent works have explored ways to revive users’ agency when building counterfactual examples to understand machine learning models, such as ViCE [11] and DECE [7]. However, to the best of our knowledge, there are no other proposals, beyond GAM Coach [33], considering both algorithmic aspects and users’ interaction aspects.

### 3 STRUCTURING INTERACTION WITH RECOUSE BASED ON PREFERENCE ELICITATION

On the side of the algorithmic recourse, our approach is inspired by the algorithm proposed by De Toni and colleagues [26] which integrates recourse with Preference Elicitation [4] to interactively estimate user preferences via pairwise comparison queries. While a detailed description and a formal evaluation of the algorithm are beyond the scope of this short paper, we present the gist of the algorithm to facilitate the understanding of the proposed *guided-interaction* paradigm. Our approach is based on the idea of eliciting user preferences via an interaction protocol that asks the user to choose the preferred intervention from a small set  $k$  of alternatives. The options are selected and proposed to the user at each algorithmic interaction. We use a greedy approach to select the alternatives maximizing the information gain [29]. The user choice is then used to update a distribution over cost weights  $P(W)$ , progressively converging towards the true user preferences  $w^*$  used by eq. 1. Lastly, this preference elicitation strategy is integrated into a reinforcement learning agent coupled with a discrete search procedure (Monte Carlo Tree Search [9]) to efficiently discover recourse solutions. This approach realizes an instance of the coactive learning strategy as proposed by [23].

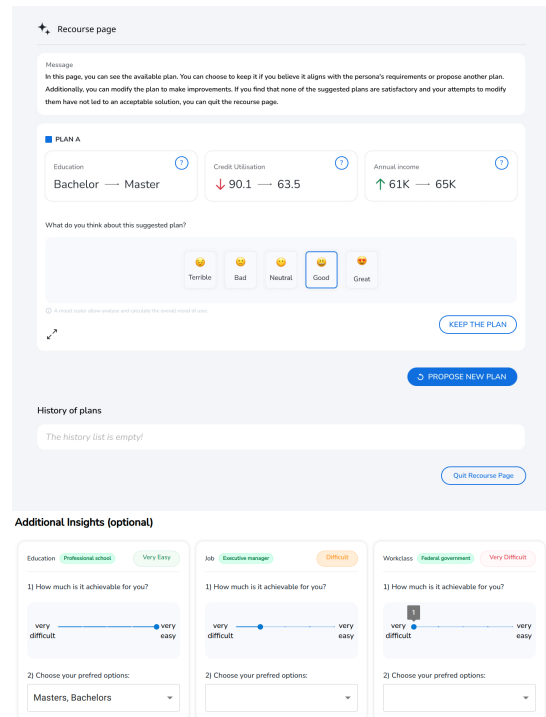
The *guided-interaction* paradigm has been designed to take advantage of this algorithmic approach (see fig. 1).

- (1) First, the user is provided an intervention that can overturn the decision given by the classifier;
- (2) Second, the user has to rate the proposed intervention with a 5-point Likert scale (“Terrible”, “Bad”, “Neutral”, “Good” and “Great”);
- (3) Then, for each feature modified by the intervention, the user can specify how achievable such change is on a 5-point Likert scale ranging from “Very difficult” to “Very easy”. Users can also specify a preferred set of options for categorical features or an acceptable range for numerical features. See fig. 1, left, for an example of the preference panel;
- (4) Lastly, the user can decide to accept the proposed intervention or ask to generate a new one.

**The control interface.** The graphical interface designed as a control condition of the study has been inspired by GAM Coach [33] but with a *look-and-feel* closer to our interface discussed above. This has been done to reduce possible biases due to non-relevant aspects. The control interface implements a *purely exploratory* perspective: one or more plans are generated, possibly considering previous constraints on features set by the user, but without learning from the user preferences. On the other hand, the interface allows the user to progressively disclose all the features and specify constraints on features not present in the suggested plan(s). Figure 2 shows the main screen of the control interface.

### 4 THE PILOT STUDY

For the pilot study, a loan application task has been chosen as a typical scenario for recourse assessment. In order to increase the number of features on which to base the decisions, two datasets:



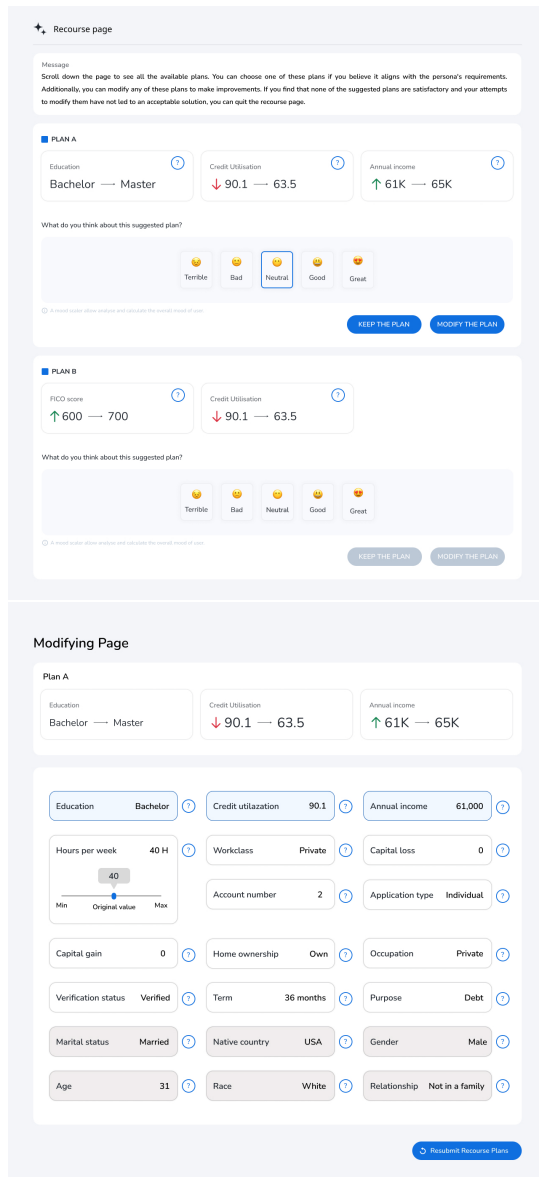
**Figure 1: The *guided* interface that uses the algorithmic recourse method enriched with preference elicitation: the user can only express preferences on the proposed features since they have been selected by the algorithm as those that optimize the learning of a personalized cost function.**

the “Adult Census Income” [1] and the “Lending Club”<sup>1</sup> have been combined. Each dataset contains both continuous and categorical features for a total of 104 attributes. We selected a subset of 31 attributes that we believe the user would be reasonably capable of acting upon (such as, Job Category and Education Level). Then, we defined an action set  $\mathcal{A}$  comprising all the features, those 31 features were marked as *actionable*, and all the others as *non-actionable* but still used in the machine learning process. As the black-box classifier  $h$ , we trained an ensemble composed of two multi-layered perceptron architectures by minimizing the empirical risk over the training data via stochastic gradient descent [20]. We trained the classifier to be as realistic as possible by using both the actionable and unactionable user attributes. In our loan scenario, we assume a loan is granted to a user only if both ensemble components give a positive classification.

**Study design.** The study employed a *within-subjects* design. Each participant was asked to solve the same task (overturn an unfavorable outcome) with the *guided*-style interface and with the *purely exploratory*-style interface as a control condition. The order of the interfaces was randomized.

**Procedure.** To make sure that the first interaction with the systems resulted in a rejection of the loan request (which was needed

<sup>1</sup><https://www.kaggle.com/datasets/wordsforthewise/lending-club>



**Figure 2: Control interface used in the study; inspired by GAM Coach [33]: the system proposes one or more plans, and the users can express their attitude and constraints on any available feature. On top, two plans each one offering the possibility of accepting it or modifying it; on bottom, one of the plans while being modified, all the actionable features can be accessed.**

to start the recourse process), we used an approach based on *personas*: fictional characters for which the specific characteristics were prepared in advanced [17]. While *personas* are commonly used by designers to describe and empathize with users, we adopted this approach in user testing by asking our users to play the role of a specific *persona* rather than responding on the interface for themselves. We also believe that this approach had the advantage of minimizing

the potential influence of previous experience, personal knowledge, and attitudes toward lending. We defined two *personas*, one female and one male. Each participant was asked to role-play one *persona* while using one interface and the other *persona* while using the other one. The association between interfaces and *personas* was also randomized.

**Measures.** We adopted a mixed methods approach employing quantitative and qualitative measures. We employed the short version of the User-Experience Questionnaire (UEQ) [21, 22]. This questionnaire consists of 26 semantic differential items on six scales: *Attractiveness*, which measures the overall impression; *Perspiciuity*, which measures the ease of use and ease of learning, as well as, perceived familiarity; *Efficiency*, how efficient and fast is perceived the product; *Dependability*, how much the user feel in control of and safe during the interaction; *Stimulation*, how exciting and motivating is to use the product; and finally *Novelty*, that measures the degree to which the product is considered innovative and creative, and how much it captures users' attention. Since the questionnaire was administered by one experimenter, we selected one time for each scale, using the following anchors: *Attractiveness*: unattractive vs attractive; *Perspiciuity*: complicated vs easy; *Efficiency*: inefficient vs efficient; *Stimulation*: boring vs exciting; *Novelty*: conventional vs inventive; *Dependability*: obstructive vs supportive. The experimenter administered the questionnaire at the end of the task with each interface. Besides scoring the semantic differential of each item, the experimenter also asked for a brief motivation. Those verbal reports were then transcribed and analyzed with Thematic Analysis [5].


**Participants.** The call for participants was promoted among local university students, and the sampling was organized through a *snowball effect*. Ultimately, 12 participants, three for each combination of *persona*-interface, were recruited. The participants had an average age of 25.6 years (standard deviation 4.3; minimum 19 and max 33 years). Among them, there were ten females and two males.

## 5 RESULTS AND DISCUSSION

Figure 3 presents the scores of the semantic differentials. The overall scores were good with an average above 4, except *Perspiciuity* and *Dependability*, whose scores ranged between 3 and 4. None of the observed differences are statistically significant, as confirmed by Friedman's ranking tests (with Bonferroni's correction).

For the *guided*-style there is a positive correlation between the duration of the task and the dimension of *Attractiveness* (Pearsons'  $r=0.26$ ), *Perspiciuity* ( $r=0.73$ ) and *Dependability* ( $r=0.15$ ) while the same correlations are negative for the *exploratory*-style interface (respectively,  $-0.23$ ,  $-0.32$ , and  $-0.12$ ); for the *Efficiency* score, the correlation is negative for the *exploratory*-style interface and not statistically different from 0 for the *guided*-style interface. *Stimulation* does not seem to correlate with duration in either one of the interfaces. While these results are preliminary due to the low number of participants, it is noteworthy that our users appreciate the freedom provided by the *explorative*-style interface, unless it leads to time wastage. Conversely, more time spent on the *guided*-style interface is perceived as useful, possibly because it offers a more direct path to finding a solution.





**Figure 3: Scores of the UEQ questionnaires for the two interfaces: the value 1 corresponds to the negative anchor and the value 5 corresponds to the positive anchor for the semantic differential items.**

From a Thematic Analysis of the verbal reports collected as motivations for the scores of the scales, it emerges that our participants articulated their experiences along two different dimensions. On one side, the freedom to propose changes and manipulate the input (even though our participants never engaged in such elaborate behaviors). On the other side, the need to avoid confusion and quickly reach a solution. For example, p1011 commenting on Perspicuity for the exploratory interface said “I could solve them without effort” and in commenting Efficiency for the guided “I barely put effort for this”. Yet, in commenting on Stimulation, she said “I was more excited to use this one [the exploratory] because I had more plans to check. I liked it more.” Nevertheless, for the guided interface, she reflected that “Maybe for someone that wants loan approval in real life, it would be more interesting.” P1013 discussing Attractiveness said “I prefer the first platform [the guided-style], in the second one [the exploratory] I could express my preferences with more details but it took too much time and confused me.” P1016 on commenting the Perspicuity of exploratory “it was harder than the first one, but this one you had more options to change I prefer this one.” She was also more explicit on commenting Stimulation for the exploratory style “[this one is] more exciting than [the one] before because I had more control over changing. The fact that you have more plans is more exciting than having one plan.” On a similar line, P018 on Stimulation of the exploratory “this one is more exciting [...] you could change the plan then you have more plans again.” P1014 on commenting on the Dependability of the exploratory summarized these terms: “I didn’t expect to receive many plans [positive surprise], and negative surprised because of a long list of information in the modifying page [the list of all the disclosed features]”. In a similar venue, P1019 on the Efficiency of the exploratory “I think it was 5, but the other one was shorter, and people (may) find the other one more efficient.”

## 6 CONCLUSION

Our work addressed a core aspect of the design of *Human-Centred AI*, namely, how to balance user control with system autonomy [25], in the context of AI-based algorithmic recourse. Our design was framed in the strategy of *prospective interaction*. Still, it differentiates from a *purely exploratory* approach by framing the exploration

within the boundaries of an optimized preference elicitation process. The aim of our work was not just to compare our interface to GAM Coach [33] but rather to compare the two approaches of *prospective* methods for AI-driven systems: a purely *explorative* approach and a purely *guided* approach. It is worth noting that they differ concerning (i) how the algorithmic recourse is realized in the backend and (ii) how the interface fosters either free interactions up to what-if explorations or rather an efficient but constrained step-by-step process. However, in both cases, users are engaged in proactively “probing the algorithm boundaries with different inputs” [25].

Our results suggest that users may recognize that the *guided* interaction paradigm improves efficiency. However, they also feel less freedom to experiment with “what-if” scenarios. Nevertheless, the time spent on the *purely exploratory* interface tends to be perceived as a lack of efficiency, which reduces attractiveness, perspicuity, and dependability. Conversely, for the *guided*-style interface, more time on the interface seems to increase its attractiveness, perspicuity, and dependability while not impacting the perceived efficiency. The evidence of this study might suggest that this type of interfaces should combine these two approaches by trying to support *explorative* behavior while gently pushing toward a *guided* effective solution. We acknowledge that the main limitations of our study include the small sample size and the likelihood that our participants may not have fully grasped the task’s difficulty due to their lack of previous experience in money lending. Nevertheless, albeit still preliminary, we believe our work may help shed new light on the ongoing debate on *Human-Centred AI* and lay the basis for further studies on this topic.

## ACKNOWLEDGMENTS

The work was partially supported by the following projects: AI@Trento, FBK-Unitn (for GdT), MUR PNRR project FAIR - Future AI Research (PE00000013) funded by the NextGenerationEU (for AP, BL,MZ), Horizon Europe Programme, grant #101120237-ELIAS (for AP, BL) and grant #101120763-TANGO (for AP, BL, KT). Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Health and Digital Executive Agency (HaDEA). Neither the European Union nor the granting authority can be held responsible for them.

## REFERENCES

- [1] Barry Becker and Ronny Kohavi. 1996. Adult. UCI Machine Learning Repository. DOI: <https://doi.org/10.24432/C5XW20>.
- [2] Siddharth Bhatore, Lalit Mohan, and Y Raghu Reddy. 2020. Machine learning techniques for credit risk evaluation: a systematic literature review. *Journal of Banking and Financial Technology* 4 (2020), 111–138.
- [3] Francesco Bodria, Fosca Giannotti, Riccardo Guidotti, Francesca Naretto, Dino Pedreschi, and Salvatore Rinzivillo. 2023. Benchmarking and survey of explanation methods for black box models. *Data Mining and Knowledge Discovery* 37, 5 (2023). <https://doi.org/10.1007/s10618-023-00933-9>
- [4] Craig Boutilier. 2002. A POMDP formulation of preference elicitation problems. In *AAAI*. 239–246.
- [5] Virginia Braun, Victoria Clarke, Nikki Hayfield, and Gareth Terry. 2019. Thematic Analysis. In *Handbook of Research Methods in Health Social Sciences*, Pranee Liamputtong (Ed.). Springer Singapore. [https://doi.org/10.1007/978-981-10-5251-4\\_103](https://doi.org/10.1007/978-981-10-5251-4_103)
- [6] Ruth MJ Byrne. 2019. Counterfactuals in Explainable Artificial Intelligence (XAI): Evidence from Human Reasoning.. In *IJCAI*. 6276–6282.

- [7] Furui Cheng, Yao Ming, and Huamin Qu. 2021. DECE: Decision Explorer with Counterfactual Explanations for Machine Learning Models. *IEEE Transactions on Visualization and Computer Graphics* 27, 2 (2021), 1438–1447. <https://doi.org/10.1109/TVCG.2020.3030342>
- [8] Giandomenico Cornacchia, Vito Walter Anelli, Giovanni Maria Biancofiore, Fedelucio Narducci, Claudio Pomo, Azzurra Ragone, and Eugenio Di Sciascio. 2023. Auditing fairness under unawareness through counterfactual reasoning. *Information Processing & Management* 60, 2 (2023), 103224. <https://doi.org/10.1016/j.ipm.2022.103224>
- [9] Rémi Coulom. 2006. Efficient selectivity and backup operators in Monte-Carlo tree search. In *International conference on computers and games*. Springer, 72–83.
- [10] Giovanni De Toni, Bruno Lepri, and Andrea Passerini. 2023. Synthesizing explainable counterfactual policies for algorithmic recourse with program synthesis. *Machine Learning* 112, 4 (2023), 1389–1409.
- [11] Oscar Gomez, Steffen Holter, Jun Yuan, and Enrico Bertini. 2020. ViCE: Visual Counterfactual Explanations for Machine Learning Models. In *Proceedings of the 25th International Conference on Intelligent User Interfaces (Cagliari, Italy) (IUI '20)*. Association for Computing Machinery, New York, NY, USA, 531–535. <https://doi.org/10.1145/3377325.3377536>
- [12] Riccardo Guidotti. 2022. Counterfactual explanations and how to find them: literature review and benchmarking. *Data Mining and Knowledge Discovery* (2022). <https://doi.org/10.1007/s10618-022-00831-6>
- [13] Amir-Hossein Karimi, Gilles Barthe, Bernhard Schölkopf, and Isabel Valera. 2020. A survey of algorithmic recourse: definitions, formulations, solutions, and prospects. *CoRR abs/2010.04050* (2020). arXiv:2010.04050 <https://arxiv.org/abs/2010.04050>
- [14] Amir-Hossein Karimi, Julius Von Kügelgen, Bernhard Schölkopf, and Isabel Valera. 2020. Algorithmic recourse under imperfect causal knowledge: a probabilistic approach. *Advances in neural information processing systems* 33 (2020), 265–277.
- [15] Ralph L. Keeney and Howard Raiffa. 1993. *Decisions with Multiple Objectives: Preferences and Value Trade-Offs*. Cambridge University Press. <https://doi.org/10.1017/CBO9781139174084>
- [16] Divyat Mahajan, Chenhao Tan, and Amit Sharma. 2019. Preserving causal constraints in counterfactual explanations for machine learning classifiers. *arXiv preprint arXiv:1912.03277* (2019).
- [17] Tara Matthews, Tejinder Judge, and Steve Whittaker. 2012. How do designers and user experience professionals actually perceive and use personas?. In *In Proceedings of CHI 2012*.
- [18] Rafael Poyiadzi, Kacper Sokol, Raul Santos-Rodriguez, Tijl De Bie, and Peter Flach. 2020. FACE: feasible and actionable counterfactual explanations. In *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*. 344–350.
- [19] G. Ramakrishnan, Y. C. Lee, and A. Albarghouthi. 2020. Synthesizing action sequences for modifying model decisions. In *AAAI*, Vol. 34. 5462–5469.
- [20] Sebastian Ruder. 2016. An overview of gradient descent optimization algorithms. *arXiv preprint arXiv:1609.04747* (2016).
- [21] Andrea Schankin, Matthias Budde, Till Riedel, and Michael Beigl. 2022. Psychometric Properties of the User Experience Questionnaire (UEQ). In *CHI Conference on Human Factors in Computing Systems*. ACM. <https://doi.org/10.1145/3491102.3502098>
- [22] Martin Schrepp, Andreas Hinderks, and Jörg Thomaschewski. 2017. Design and Evaluation of a Short Version of the User Experience Questionnaire (UEQ-S). *International Journal of Interactive Multimedia and Artificial Intelligence* 4, 6 (2017). <https://doi.org/10.9781/ijimai.2017.09.001>
- [23] Pannaga Shivaswamy and Thorsten Joachims. 2015. Coactive Learning. *J. Artif. Intell. Res. (JAIR)* 53 (2015), 1–40. <https://doi.org/10.1613/jair.4539>
- [24] Ben Shneiderman. 2020. Bridging the Gap Between Ethics and Practice: Guidelines for Reliable, Safe, and Trustworthy Human-centered AI Systems. *ACM Transactions on Interactive Intelligent Systems* 10, 4 (Dec. 2020), 1–31. <https://doi.org/10.1145/3419764>
- [25] Ben Shneiderman. 2022. *Human-centered AI*. Oxford University Press.
- [26] Giovanni De Toni, Paolo Viappiani, Stefano Teso, Bruno Lepri, and Andrea Passerini. 2024. Personalized Algorithmic Recourse with Preference Elicitation. *Transactions on Machine Learning Research* (2024). <https://openreview.net/forum?id=8sg2I9zXgO>
- [27] Berk Ustun, Alexander Spangher, and Yang Liu. 2019. Actionable Recourse in Linear Classification. In *Proceedings of the Conference on Fairness, Accountability, and Transparency*. ACM. <https://doi.org/10.1145/3287560.3287566>
- [28] Sahil Verma, Varich Boonsanong, Minh Hoang, Keegan E Hines, John P Dickerson, and Chirag Shah. 2020. Counterfactual explanations and algorithmic recourses for machine learning: A review. *arXiv preprint arXiv:2010.10596* (2020).
- [29] Paolo Viappiani and Craig Boutilier. 2010. Optimal bayesian recommendation sets and myopically optimal choice query sets. *Advances in neural information processing systems* 23 (2010).
- [30] Paul Voigt and Axel Von Dem Bussche. 2017. *The EU General Data Protection Regulation (GDPR)*. Springer International Publishing, Cham. <https://doi.org/10.1007/978-3-319-57959-7>
- [31] Sandra Wachter, Brent Mittelstadt, and Chris Russell. 2017. Counterfactual Explanations Without Opening the Black Box: Automated Decisions and the GDPR. *SSRN Electronic Journal* (2017). <https://doi.org/10.2139/ssrn.3063289>
- [32] Caroline Wang, Bin Han, Bhrij Patel, and Cynthia Rudin. 2023. In pursuit of interpretable, fair and accurate machine learning for criminal recidivism prediction. *Journal of Quantitative Criminology* 39, 2 (2023), 519–581.
- [33] Zijie J Wang, Jennifer Wortman Vaughan, Rich Caruana, and Duen Horng Chau. 2023. GAM Coach: Towards Interactive and User-centered Algorithmic Recourse. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems (CHI '23)*. ACM.