

# ReVise: A Human-AI Interface for Incremental Algorithmic Recourse

Kaustav Bhattacharjee\*  
Optimization & Control Group  
Pacific Northwest National Laboratory

Jun Yuan†  
Department of Data Science  
New Jersey Institute of Technology

Aritra Dasgupta‡  
Department of Data Science  
New Jersey Institute of Technology

## ABSTRACT

The recent adoption of artificial intelligence in socio-technical systems raises concerns about the black-box nature of the resulting decisions in fields such as hiring, finance, admissions, etc. If data subjects—such as job applicants, loan applicants, and students—receive an unfavorable outcome, they may be interested in algorithmic recourse, which involves updating certain features to yield a more favorable result when re-evaluated by algorithmic decision-making. Unfortunately, when individuals do not fully understand the incremental steps needed to change their circumstances, they risk following misguided paths that can lead to significant, long-term adverse consequences. Existing recourse approaches focus exclusively on the *final* recourse goal but neglect the possible incremental steps to reach the goal with real-life constraints, user preferences, and model artifacts. To address this gap, we formulate a visual analytic workflow for *incremental* recourse planning in collaboration with AI/ML experts and contribute an interactive visualization interface that helps data subjects efficiently navigate the recourse alternatives and make an informed decision. We present a usage scenario and subjective feedback from observational studies with twelve graduate students using a real-world dataset, which demonstrates that our approach can be instrumental for data subjects in choosing a suitable recourse path.

**Index Terms:** Visual analytics, recourse, machine learning.

## 1 INTRODUCTION

Artificial intelligence (AI) drives algorithmic decision-making in various socio-technical contexts such as hiring, lending, and admissions, where individuals — serving as data subjects — experience tangible impacts based on the outcomes produced by these systems. Data subjects who receive undesired outcomes may seek algorithmic recourse, i.e., the process of updating certain features to yield improved outcomes when re-evaluated by the decision-making process [10, 18]. Yet, achieving effective recourse is complex; it must account for the data subject’s initial feature values, their unique preferences and real-life constraints, the decision-making rules of the model, and the influence of other data subjects processed by the model. Two main forms of recourse exist in the literature: contrastive explanation [17], which offers examples of others who received favorable outcomes, and consequential recommendations [22], which suggest actions to achieve favorable outcomes in the future. However, both approaches focus predominantly on the final outcome rather than on the step-by-step process of reaching recourse. This emphasis on the ultimate goal, rather than on the incremental journey, poses challenges in practical, real-world scenarios where gradual adjustments are often necessary.

Our review of the literature on algorithmic recourse reveals two principal limitations in current approaches. First, many methods treat recourse as a deterministic problem aimed solely at achieving

a specific favorable outcome, often neglecting the incremental progression from a negative state. For example, a loan applicant might be instructed to boost their credit score by a fixed amount to secure approval, without considering that incremental improvements could gradually enhance their likelihood of success. Second, existing frameworks tend to prescribe changes in feature values without allowing individuals the flexibility to select tailored actions along their improvement journey. One might imagine a job applicant being advised to pursue a single, rigid upgrade—such as obtaining another degree—instead of exploring a series of manageable steps that align with their personal constraints and preferences. This narrow focus on binary endpoints disregards the inherent uncertainties of real-world settings, where understanding gradual paths is critical. To address these problems, we introduce *incremental recourse*, guiding data subjects through a sequence of steps toward desired outcomes rather than presenting a single change. Instead of focusing on binary outcomes, we utilize probability scores to show gradual progress—for example, using the probability of loan approval rather than simply approval/disapproval designations.

We implement a human-in-the-loop approach for incremental recourse planning where users actively participate in defining their recourse path. Users update their data to define subsequent recourse states based on their preferences, real-life constraints, and insights from other data subjects’ experiences. This approach requires understanding model behavior locally, which presents challenges when working with black-box models. We incorporate feature attribution explanations in our interactive workflow as they represent local model preferences, allowing users to make informed choices about their next steps rather than randomly testing for better outcomes. We contribute *ReVise*, an interactive recourse planning assistant, which facilitates multi-way comparison and decision-making among recourse states and paths, supported by interactive visualization designs for human-in-the-loop recourse calibration. We focus specifically on decision-making tasks that underpin the visualization design space, moving beyond the sense-making or exploration tasks commonly addressed in visualization research [4]. Our visual analytic approach enables users to compare alternative paths efficiently, understand trade-offs between different feature changes, and make informed decisions about their recourse journey with greater agency and awareness of incremental probability improvements.

## 2 RELATED WORK

Algorithmic recourse [11] generally follows two main approaches: contrastive explanation [17] (“applicants with credit scores above 720 were approved”) and consequential recommendations [22] (“increase your credit score by 50 points to improve approval chances”). Our focus on *incremental* recourse differs from both by emphasizing momentum toward favorable outcomes rather than guaranteeing them, constructing action sequences where features may be improved multiple times. Karimi et al. argue that structured causal models are essential for improving the feasibility of recourse, distinguishing between mere explanations and meaningful interventions for actionable change [12]. However, developing precise causal models is challenging and often influenced by individual perspectives, which can lead automated recourse to produce unrealistic recommendations. In this work, we offer visual inter-

\*e-mail: kaustav.bhattacharjee@pnnl.gov

†e-mail: jy448@njit.edu

‡e-mail: aritra.dasgupta@njit.edu

pretations of actionable incremental targets, enabling users to adjust and customize the machine-generated suggestions according to their own preferences.

Feature attribution methods quantify the contribution of each feature to a model’s prediction, an insight vital for algorithmic recourse as it highlights which feature modifications most effectively change outcomes. Researchers have linked feature attribution with recourse via counterfactual explanations—unifying attribution and counterfactuals [14] and creating counterfactual Shapley additive explanations [1]. Additional approaches have leveraged techniques such as: generating counterfactual and contrastive explanations using SHAP [20]; employing Shapley values to estimate the combinatorial effect of multiple changes to generate counterfactual text [5]; and proposing methods to derive counterfactual explanations using feature attributions generated by LIME and SHAP [19]. However, some critiques argue that using attribution-based explanations as final recourse suggestions may not be robust enough [7] or require reformulation for greater intuitiveness [6]. In this work, we utilize feature attribution to form visual cues for incremental counterfactual suggestions by limiting explanations to local regions, which facilitates a practical assessment of actionability and forms the foundation of *ReVise*.

Existing visualization-based interfaces for recourse include DECE, which uses contrastive examples within user-defined feature ranges for neural network predictions [2]; GAM Coach’s text-heavy, tab-based interface [23]; and tools such as the what-if tool [24] and ViCE [9]. GAM Coach’s reliance on tabbed navigation hinders multi-way comparisons and offers limited support for visualizing progressive recourse trajectories, and none of these systems explicitly facilitate decision-making tasks—investigating, synthesizing, and selecting among alternatives—which are crucial for incremental recourse planning [4]. To the best of our knowledge, our work is the first in visual analytics to bridge the gap between incremental recourse needs, the corresponding metrics, and the interactive visualizations required to effectively communicate them to end users.

### 3 VISUAL ANALYTIC TASKS & INTERFACE DESIGN

The design space for developing a visualization system for incremental recourse planning is driven by the need for a transparent explanation of the recourse steps and the clear evaluation of potential next steps. In this section, we first discuss the visual analytic tasks we formulated to achieve these goals in our web-based interface, and then explore the supporting design elements.

#### 3.1 Visual Analytic Tasks

We have identified six key visual analytic tasks that form the basis of the design space for incremental recourse planning—the first three support understanding the current state, while the next three facilitate informed decision-making for the next recourse state.

**T1: Understanding current state features and attributions:** The first step in recourse planning is to understand the current state across various features and their corresponding attributions. This task relates to the understanding and comparison of the feature values for the current state and their attributions.

**T2: Assessing impact of features on model outcomes:** Understanding the model outcome and the impact of each feature on the outcome is an essential step for recourse planning. This task addresses understanding the model outcome in relation to each feature.

**T3: Comparing feature deviation from current state averages:** A user may wish to understand how their current feature values deviate from averages within normalized feature ranges, enabling them to identify which features are above or below typical values and prioritize improvements accordingly.

**T4: Identifying potential targets for next state:** Users may be willing to understand which targets would be suitable for the next state

based on the projection values of all the possible targets. This task revolves around aiding the users to find a target for the next state from a pool of candidate targets.

**T5: Comparing potential targets across features:** The potential targets may have different attributions for each feature. Hence, it is essential to compare them across different features to refine the selection of the potential targets.

**T6: Comparing recourse paths:** Selecting a target for the next state creates a recourse path, which can be incrementally modified to reach the desired target. This task relates to the iterative modification of the recourse paths by comparing the model outcomes and feature attributions through different recourse states.

#### 3.2 Visual Cues for Expressing Recourse Patterns

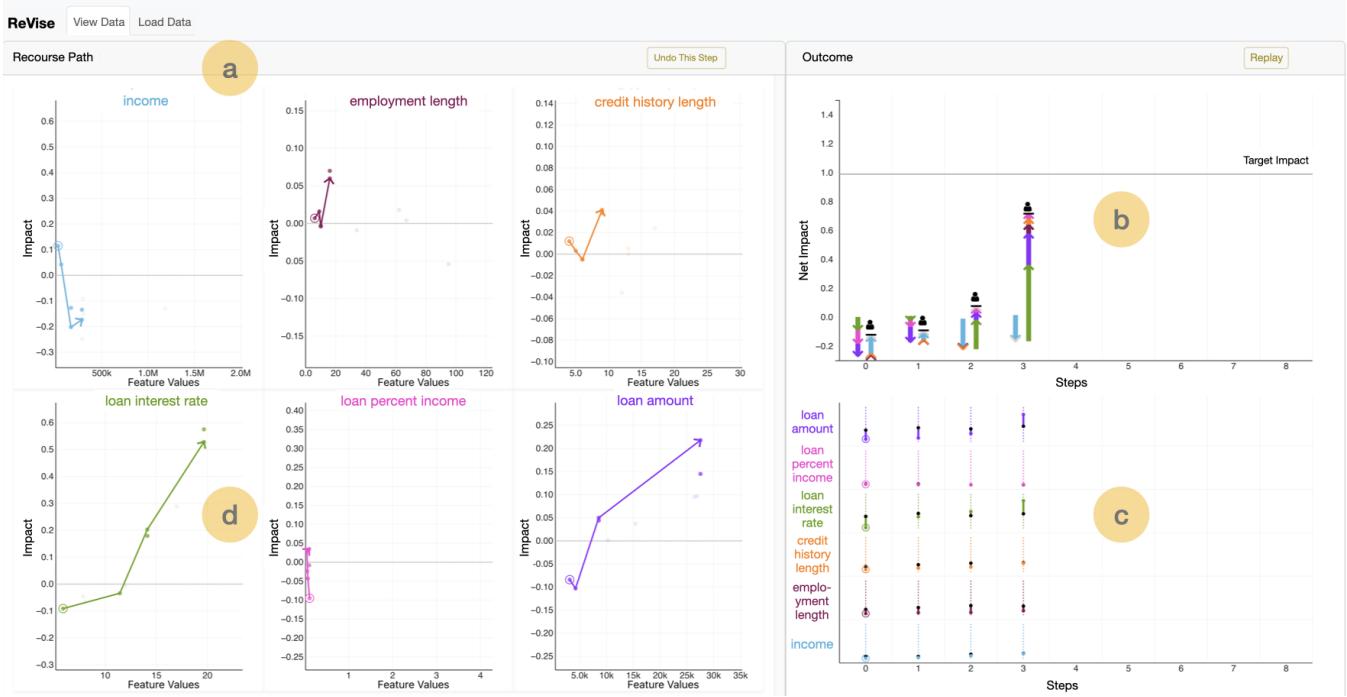
In this subsection, we align recourse components with design elements, highlighting the cues and visual analytic tasks that facilitate incremental recourse for one or more data subjects.

**Cues for Associating Feature Attribution and Value:** The attribution for a data subject’s feature is determined and estimated through a feature explanation method (e.g., SHAP [16]) which distributes the outcome across all features of the data subject. Each data subject in the current state is depicted as a dot in a scatterplot (Figure 1a), where the dot’s position reflects its feature value and corresponding attribution—for instance, a loan applicant’s employment length of 5 years might have a positive attribution of 0.15 toward loan approval (T1). Hovering over a dot highlights its position relative to the overall current state, enabling comparison of potential targets for the next state (T5).

**Cues for Part-to-Whole Impact:** Different features have varying impacts on the model outcome for a particular data subject at a certain state. The impact of each feature is encoded by the length of its attribution, also illustrating its contribution relative to the overall outcome (Figure 1b) (T2). For example, this design may reveal that the loan interest rate has a more significant effect than the loan amount for a particular data subject, suggesting that improving the interest rate could yield better outcomes in the next recourse step. To assess feature value deviations at each recourse step (T3), we represent each feature using a dotted line that spans the normalized range, with a black dot indicating the average feature value across all available recourse states and a colored dot representing the current feature value (Figure 1c). A solid colored line connects these dots to visualize the deviation—whether an increase or decrease—from the average, enabling users to quickly identify features that differ significantly from the norm for strategic decision-making about which features to prioritize in the next recourse state.

**Cues for Direction:** Another important aspect of understanding a feature’s impact on the model outcome is the direction of its attribution, as a data subject may find all features contributing negatively in one state, while in another, some features may contribute positively and others negatively (Figure 1b) (T2). By addressing features with negative impacts—selecting a potential target for the next state with higher attribution for those features—a more favorable recourse path may be achieved. The direction of the recourse trajectory indicates whether the chosen path is leading toward improvement; an upward trajectory signifies progress (Figure 1d), whereas a downward trend suggests that the selected path may be suboptimal. This encoding allows the data subject to backtrack to a previous state and iterate through a different recourse path based on alternative potential targets (T6).

**Cues for Actionability:** Selecting the potential target for the next state can be challenging given the pool of candidate targets with varying feature values and attributions. To overcome this, we use a change-to-attribution ratio (also called projection) to identify the top targets for the next state (T4), visually encoding it through dot size with the top three targets highlighted by larger dots. We considered an alternative design that resized all dots based on their pro-



**Figure 1: Recourse Planning Usage Scenario:** (a) Utilizing the Credit Risk dataset, the user selected a recourse path from the Recourse Path Explorer. (b) After multiple iterations, she identified a recourse path that resulted in significant improvement in the model outcome. (c) At the same time, the user observed that some features reached significantly higher-than-average values, while others showed a downward trend in their attributions. (d) Despite these differences, the majority of features followed an upward trajectory, consistent with the overall aim of recourse planning. Ultimately, the user confirmed that *ReVise*'s visual encodings faithfully reflect the explanation results offered by SHAP.

jection values, but feedback showed that abstracting to only the top few targets was more effective.

**Cues for Trajectory:** The slope of the recourse path indicates the change magnitude; a steeper slope reflects a larger attribution gain per unit feature change, aligning with the “minimum effort, maximum output” principle—as when a two-year increase in credit history length yields a 70% attribution rise. Users can iterate through different paths to find steeper, more promising trajectories (T6).

### 3.3 Interactive Interface

In this subsection, we introduce the web-based interface for *ReVise*, designed to facilitate incremental recourse planning for data subjects and collaborating explainable AI experts. The interface comprises two main views: the Recourse Path Explorer (Figure 1a) assists users in initiating recourse planning, understanding their current state, and identifying potential targets (T1, T4, T5, T6), while the Outcome Monitor (Figures 1b and 1c) tracks model outcomes with corresponding feature impacts (T2) and shows deviations from average feature values (T3). Developed with React.js and D3.js, this coordinated layout streamlines the exploration of recourse paths while simultaneously assessing their effectiveness.

#### 3.3.1 Coordinated Multiple Views

**Recourse Path Explorer:** The Recourse Path Explorer presents six scatter plots arranged in two rows to visualize correlations between feature values and attributions for the most influential features, with each dot representing a data subject (Figure 1a). Users can select any dot, causing corresponding dots across all plots to be highlighted; selecting multiple dots creates solid lines connecting them in sequence, with arrows indicating the trajectory of the formed recourse path (Figure 1d). This interactive process aids users in understanding the current state across all features (T1).

In this view, hovering over dots enables direct comparison with the last selected point, supporting comparisons across multiple views to help users determine the next state (T5) and form effective recourse paths (T6). For example, when a user selects their data point on the Recourse Path Explorer, *ReVise* instantly displays the associated outcome and feature attributions, with large bubbles indicating actionable next steps (T4); hovering over these bubbles further highlights differences in feature values and attributions between the current and potential next states (T1, T5).

**Outcome Monitor:** Once a dot is selected on the Recourse Path Explorer, the Outcome Monitor displays a modified connected scatter plot paired with a modified deviation plot. In the upper half design, the x-axis represents the recourse path steps while the y-axis depicts the model outcome and its feature attributions (Figure 1b). For each state, negative attributions are stacked downward from a zero baseline, and positive attributions are built upward from their lowest point, with the endpoint marked by the human icon showing the resulting model outcome. The lower half maintains recourse steps on the x-axis with individual features on the y-axis (Figure 1c), using dotted lines to represent normalized feature ranges, black dots for average values, and colored dots for current values—connected by solid lines to visualize deviations from averages. Repeated for each recourse state, this visual analytic design allows users to gauge features’ impact on outcomes (T2) and assess goal proximity by comparing the positive attribution stack endpoints with target outcomes. The modified deviation plot further reveals which features lie significantly above or below the norm—indicating areas for refinement (T3)—and together, these visualizations provide an overall view of progress along the recourse path and the adjustments needed to meet the target outcome.

#### 3.3.2 Color and Interactivity

*ReVise* uses a discrete set of colors to visualize six key features and

the recourse path, following explainable AI practices while grouping the rest as “others.” Since SHAP is additive, this subset approach doesn’t violate constraints and reduces cognitive load while allowing users to select any features for path construction. Though features beyond the top six are excluded from the Recourse Path Explorer, their collective influence appears in gray on the Outcome Monitor, ensuring awareness of their impact. When users select a dot, *ReVise* initiates a sequence of data operations where a Flask-based backend identifies potential targets as per defined constraints, recalculates SHAP attributions, and returns data to adjust dot coordinates with animated transitions as dots reposition and resize. Through repeated selections, a recourse path forms across views with special encoding for starting points (unfilled circles) and endpoints (filled circles), while deselection enables backtracking. This coordinated interactivity between views facilitates effective incremental planning by allowing exploration of alternative paths.

#### 4 USAGE SCENARIO: CREDIT RISK REOURSE PLANNING

We present a usage scenario derived from pilot studies to demonstrate how visual analytics could support incremental recourse planning. Let us consider a scenario where, Allen, a business professional, is interested in understanding her chances for credit approval. Allen began by selecting the Credit Risk dataset [3] and examining feature values and attributions (**T1**). She identified a dot representing income at \$70k with higher attribution—a counterintuitive finding since lower income generally correlates with lower loan probability (Figure 1a). In the Outcome Monitor, she saw that despite positive income attribution, other features negatively influenced the loan decision (**T2**) (Figure 1b). In the lower half, she noted that most features were near average, but loan interest rate and loan amount were well below average (**T3**) (Figure 1c). This aligns with expectations that lower income combined with less favorable interest rates reduces approval chances, as higher-income applicants typically secure better loan terms.

After achieving the initial goal of understanding the current state through feature attributions and the model outcome, Allen proceeded to identify potential targets for the next state. *ReVise* automatically suggested three potential targets, depicted as larger dots (**T4**), and she hovered over these and other candidates to assess their distance from the current state and check for increased attribution across features (**T5**), paying close attention to steeper slopes indicating small feature changes yielding large positive attributions given their real-life impact. When a selected state didn’t significantly improve the model outcome, Allen used the ‘Undo’ button to revert; through several iterations, she eventually identified a recourse path yielding substantial growth in the model outcome (~80% approval odds) (Figure 1b). Thus, *ReVise* facilitated comparing different recourse paths, enabling selection of one with maximum outcome growth and minimal changes to features like income, employment length, and credit history length—which are generally within an applicant’s capacity (**T6**).

While analyzing the recourse path, Allen noted that certain features like income and employment length exhibited partial downward trajectories despite outcome improvements, but loan interest rate, loan amount, and credit history length displayed upward trajectories (Figure 1d)-contributing to high model outcomes (Figure 1b). Although credit history length and employment length showed large attribution increases from the initial state, their positive impact was smaller than loan interest rate and loan amount in the final recourse state (Figure 1b), suggesting feature importance varies across different recourse stages. The significant outcome improvement was driven by increased loan interest rates and decreased loan percent income, aligning with model explanations and confirming *ReVise*’s visual encodings faithfully represent these relationships.

#### 5 SUBJECTIVE USER FEEDBACK

To evaluate the faithfulness of *ReVise*’s visual encodings to SHAP attributions and its effectiveness in guiding recourse path generation for users, we conducted observational studies with 12 graduate students who had a background in computation and varied levels of familiarity with explainable AI techniques. Because of the inherent individual differences in user preferences for recourse planning [15], designing metrics for evaluating accuracy or efficiency is challenging. Therefore, the scope of our evaluation was anchored on assessing the subjective user experience [13] with a focus on observing if *ReVise* users can reliably engage in recourse-relevant decision-making tasks. Each session lasted approximately 45-60 minutes, during which participants completed recourse planning tasks while verbalizing their thought processes. From the feedback collected, we identified several key themes: **i) Impact of Visualizations:** Participants consistently highlighted the system’s ability to reveal feature relationships and their impact on outcomes, demonstrating how visual analytics can bridge the gap between complex model behavior and human-interpretable patterns. One participant noted, “*The interface is very useful in seeing the impact of different attributes on an end result. It was very well-organized and interactions were clear,*” while another commented, “*Very neat. Provides a fast visualization to help with planning.*” The visual encodings were particularly appreciated, with one participant stating, “*Visualizations are interactive and well represented.*” **ii) Easy to Use:** Participants recognized the system’s adaptability across different recourse scenarios, reinforcing the importance of accommodating diverse analytical approaches in visual analytics systems. “*Cool interface for analyzing/recoursing for any purpose in general,*” observed one participant, while another remarked, “*It’s visually appealing, and I like that the results are based on real students’ data,*” **iii) Advantages over existing approaches:** Participants who had prior experience with making recourse decisions appreciated the incremental flavor of recourse planning in *ReVise* as compared to traditional trial-and-error approaches. One mentioned, “*I found the ReVise interface to be very interesting... I did see it as very valuable for helping someone make choices regarding how to improve aspects such as grades or GPA*”- highlighting the value of data-driven recommendations over guesswork. **iv) Scope for Improvement:** The primary critique centered on the initial learning curve. One participant shared, “*At first glance, it may not look useful, and I felt a bit confused about how it was going to help me.*” Another noted difficulty in understanding “*how the features correlated and change,*” attributing this to “*having little knowledge about the background.*” However, several participants also highlighted the value of embedded guidance, with one noting, “*After I read the helper texts, everything became more clear, and I found those instructions the most useful ones.*”

#### 6 CONCLUSION

In this work, we propose *ReVise*, an interactive visual analytic interface for incremental recourse planning that helps users navigate and select effective recourse paths to achieve favorable outcomes from black-box machine learning models. This work advances the state of algorithmic recourse by shifting the focus from deterministic outcomes to probability-aware incremental paths, thereby providing users with greater agency when navigating the increasingly complex algorithmic systems that influence their lives. For future work, we will extend the interface to support recourse planning for groups with different strategies [8], handle incremental model shifts, analyze recourse robustness [21], enable flexible feature selection and projection, elicit user preferences, and incorporate time-span constraints for recourse steps.

## SUPPLEMENTAL MATERIALS

A demonstration video is available at <https://youtu.be/PkpE7ZHocQ>.

## ACKNOWLEDGMENTS

This work was funded in part by the National Science Foundation (NSF) grants 2312932 and 2326195. This paper is released as PNNL-SA-212903.

## REFERENCES

- [1] E. Albini, J. Long, D. Dervovic, and D. Magazzeni. Counterfactual Shapley Additive Explanations. In *2022 ACM Conference on Fairness, Accountability, and Transparency*, pp. 1054–1070. ACM, Seoul Republic of Korea, June 2022. doi: 10.1145/3531146.3533168 2
- [2] F. Cheng, Y. Ming, and H. Qu. DECE: Decision explorer with counterfactual explanations for machine learning models. *IEEE Transactions on Visualization and Computer Graphics*, 27(2):1438–1447, 2020. 2
- [3] Credit Risk Dataset. <https://www.kaggle.com/datasets/laotse/credit-risk-dataset/data>. Accessed: 2024-03-31. 4
- [4] E. Dimara and J. Stasko. A Critical Reflection on Visualization Research: Where Do Decision Making Tasks Hide? *IEEE Transactions on Visualization and Computer Graphics*, 28(1):1128–1138, Jan. 2022. doi: 10.1109/TVCG.2021.3114813 1, 2
- [5] X. Fern and Q. Pope. Text Counterfactuals via Latent Optimization and Shapley-Guided Search. In M.-F. Moens, X. Huang, L. Specia, and S. W.-t. Yih, eds., *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pp. 5578–5593. Association for Computational Linguistics, Online and Punta Cana, Dominican Republic, Nov. 2021. doi: 10.18653/v1/2021.emnlp-main.452 2
- [6] C. Fernández-Loría, F. Provost, and X. Han. Explaining Data-Driven Decisions made by AI Systems: The Counterfactual Approach, Oct. 2021. 2
- [7] H. Fokkema, R. de Heide, and T. van Erven. Attribution-based Explanations that Provide Recourse Cannot be Robust. *Journal of Machine Learning Research*, 24(360):1–37, 2023. 2
- [8] J. Fonseca, A. Bell, C. Abrate, F. Bonchi, and J. Stoyanovich. Setting the right expectations: Algorithmic recourse over time. In *Proceedings of the 3rd ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization*, pp. 1–11, 2023. 4
- [9] O. Gomez, S. Holter, J. Yuan, and E. Bertini. ViCE: Visual counterfactual explanations for machine learning models. In *Proceedings of the 25th International Conference on Intelligent User Interfaces*, pp. 531–535. ACM, Cagliari Italy, Mar. 2020. doi: 10.1145/3377325.3377536 2
- [10] R. Guo, L. Cheng, J. Li, P. R. Hahn, and H. Liu. A survey of learning causality with data: Problems and methods. *ACM Computing Surveys (CSUR)*, 53(4):1–37, 2020. 1
- [11] A.-H. Karimi, G. Barthe, B. Schölkopf, and I. Valera. A Survey of Algorithmic Recourse: Contrastive Explanations and Consequential Recommendations. *ACM Computing Surveys*, 55(5):1–29, May 2023. doi: 10.1145/3527848 1
- [12] A.-H. Karimi, B. Schölkopf, and I. Valera. Algorithmic Recourse: From Counterfactual Explanations to Interventions. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, pp. 353–362. ACM, Virtual Event Canada, Mar. 2021. doi: 10.1145/3442188.3445899 1
- [13] S. Koh, B. H. Kim, and S. Jo. Understanding the user perception and experience of interactive algorithmic recourse customization. *ACM Transactions on Computer-Human Interaction*, 31(3):1–25, 2025. 4
- [14] R. Kommiya Mothilal, D. Mahajan, C. Tan, and A. Sharma. Towards Unifying Feature Attribution and Counterfactual Explanations: Different Means to the Same End. In *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society*, pp. 652–663. ACM, Virtual Event USA, July 2021. doi: 10.1145/3461702.3462597 2
- [15] H. Lam, E. Bertini, P. Isenberg, C. Plaisant, and S. Carpendale. Empirical studies in information visualization: Seven scenarios. *IEEE transactions on visualization and computer graphics*, 18(9):1520–1536, 2011. 4
- [16] S. M. Lundberg and S.-I. Lee. A unified approach to interpreting model predictions. *Advances in neural information processing systems*, 30, 2017. 2
- [17] T. Miller. Contrastive Explanation: A Structural-Model Approach. *The Knowledge Engineering Review*, 36:e14, 2021. doi: 10.1017/S0269888921000102 1
- [18] R. Poyiadzi, K. Sokol, R. Santos-Rodriguez, T. De Bie, and P. Flach. FACE: feasible and actionable counterfactual explanations. In *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*, pp. 344–350, 2020. 1
- [19] Y. Ramon, D. Martens, F. Provost, and T. Evgeniou. A comparison of instance-level counterfactual explanation algorithms for behavioral and textual data: SEDC, LIME-C and SHAP-C. *Advances in Data Analysis and Classification*, Sept. 2020. doi: 10.1007/s11634-020-00418-3 2
- [20] S. Rathi. Generating Counterfactual and Contrastive Explanations using SHAP, June 2019. 2
- [21] S. Upadhyay, S. Joshi, and H. Lakkaraju. Towards Robust and Reliable Algorithmic Recourse. In *Advances in Neural Information Processing Systems*, vol. 34, pp. 16926–16937. Curran Associates, Inc., 2021. 4
- [22] B. Ustun, A. Spangher, and Y. Liu. Actionable Recourse in Linear Classification. In *Proceedings of the Conference on Fairness, Accountability, and Transparency*, pp. 10–19. ACM, Atlanta GA USA, Jan. 2019. doi: 10.1145/3287560.3287566 1
- [23] Z. J. Wang, J. Wortman Vaughan, R. Caruana, and D. H. Chau. GAM Coach: Towards interactive and user-centered algorithmic recourse. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, pp. 1–20, 2023. 2
- [24] J. Wexler, M. Pushkarna, T. Bolukbasi, M. Wattenberg, F. Viegas, and J. Wilson. The What-If Tool: Interactive Probing of Machine Learning Models. *IEEE Transactions on Visualization and Computer Graphics*, pp. 1–1, 2019. doi: 10.1109/TVCG.2019.2934619 2