FROM COUNTERFACTUAL EXPLANATIONS TO CAUSAL ALGORITHMIC RECOURSE: A REVIEW OF METHODS UNDER IMPERFECT CAUSAL KNOWLEDGE .

Abstract

Counterfactual explanations and algorithmic recourse aim to tell individuals how to change an unfavourable automated decision, but recommendations are useful only if they are feasible, actionable, and causally valid. This review examines how recent methods meet these requirements and where they fall short. The literature was searched mainly in Google Scholar for papers published between 2018 and 2026, and 20 studies on causal counterfactual explanations and algorithmic recourse were reviewed. The analysis identified three gaps. First, under real-world constraints, methods such as DiCE and FACE produce diverse, feasible, and actionable counterfactuals, but without modelling causal relationships, and later work shows that recourse can be fragile or invalidated over time. Second, many methods depend on a fully known structural causal model (SCM) or causal graph; recent work relaxes this through probabilistic, learned, and graph-free approaches, yet these still rely on assumptions about the available causal information. Third, hidden confounding and non-identifiability make causal effects hard to estimate; bounding approaches and multi-dataset methods address this, but depend on known confounding structure or compatible data, and learned SCMs offer limited guarantees of reliable counterfactual inference. Overall, the field has moved from single counterfactual explanations toward methods that operate under partial causal knowledge, yet no reviewed method resolves all three gaps at once. Future research should address feasibility, confounding, and incomplete causal knowledge together.

Key words:-Counterfactual Explanations, Algorithmic Recourse, Structural Causal Models, Imperfect Causal Knowledge, Causal Inference, Non-Identifiability.

Introduction:-

Artificial intelligence (AI) is widely used in important decision-making processes across various industries, including healthcare, education, and insurance. Given the impact of these decisions, it is important to understand how they are made; this is the focus of eXplainable Artificial Intelligence (XAI), a field concerned with understanding why an AI model makes a specific decision. Several methods fall under XAI, including SHapley Additive exPlanations (SHAP) and Local Interpretable Model-agnostic Explanations (LIME); this review, however, focuses on a third approach: counterfactuals. What sets counterfactuals apart from these other methods is that they do not simply justify the AI's decision; instead, they provide a predicted way of changing the decision made by the AI. For example, if a person applies for a loan and is rejected by a bank's ML model, the counterfactual explanation would not explain why the applicant was rejected, but would instead state something like, "If the applicant had reduced their existing

UNDER PEER REVIEW IN IJAR

credit card debt by $2,000, the loan would have been approved" [1], [2], helping the person understand how a specific outcome could be changed.

One of the major challenges in the generation of counterfactuals is that, oftentimes, these counterfactuals are not actionable or feasible; that is, the counterfactual provides a method of changing the outcome that proves to be impractical for the user. This review addresses studies that propose methods to address this problem using causal approaches, in order to align recourse with practical settings and make it realistic [3], [4], [5]. The problem with such methods is that they typically depend heavily on complete knowledge of the underlying structural causal model (SCM), complete causal graphs, and causal relationships, which are, in most cases, unavailable in real world scenarios - this renders such methods impractical under realistic conditions [3], [6], [7]. This review also examines studies that relax the assumption of such ideal conditions and operate with incomplete or no causal knowledge. Another challenge that disrupts the generation of causal counterfactuals is "hidden confounding," which occurs when an unobserved or hidden variable influences the outcome of the algorithmic recourse [8]. This review therefore also examines studies focused on addressing this problem, as well as studies targeting the scalability, adversarial robustness, and temporal stability of recourse recommendations [9], [10], [11].

Overall, this review targets three main research gaps:

1) **Feasibility and actionability under real-world constraints:**

   A counterfactual may be valid but still impossible or unrealistic for a person to act on.

2) **Dependence on a fully known SCM/causal graph:**

   Most causal recourse methods assume the true SCM or causal graph is completely known, which is rarely true in real-world settings.

3) **Confounding and non-identifiability:**

   Even with a known causal graph, unobserved variables (confounders) or multiple SCMs fitting the same data can make counterfactual predictions unreliable.

The paper is organized as follows:

Section II presents the research methodology, Section III reviews the relevant literature, Section IV discusses the findings and future research directions, and Section V concludes the review.

RESEARCH METHODOLOGY

UNDER PEER REVIEW IN IJAR

The literature reviewed for this research was primarily identified using Google Scholar between July and August 2026. The search terms used to find and select the papers varied throughout the process. The search began with "counterfactuals in XAI" as a broad topic, without any restriction on publication date, before narrowing to a causal direction with specific constraints such as unknown SCM, using search terms like "algorithmic recourse under causal uncertainty" and "causal counterfactuals with incomplete structural causal models." The research was also narrowed down to research papers between 2018 and 2026 to make sure that the research is relevant to today's time as the field is evolving rapidly.

The earlier studies reviewed were relatively broad in scope, covering the general topic of counterfactuals. The review was then narrowed to include only papers that targeted fixing a specific problem related to one of the three identified research gaps. Studies were excluded if they focused on any direction other than causality, or causality under specific constraints and limitations. The initial searches, conducted using various search terms, yielded approximately 40-50 candidate papers, of which 20 were ultimately selected, studied, and summarized.

For each included study, the review focused on the problem it identifies and presents, how it aims to solve that problem and to what extent, the constraints under which the experiments were conducted or the method was developed (e.g., unknown SCM, confounding variables), the limitations of the method, and the future work suggested by the authors. Studies with similar limitations were then grouped together to develop a clearer conceptual understanding of the common limitations they shared, making it easier to identify the recurring gaps.

LITERATURE REVIEW

Table I. Reviewed studies grouped by research gap.

S.NO	IDENTIFIED RESEARCH GAP	Papers
1.	Feasibility and actionability under real-world constraints	[2],[4],[9],[10],[11]
2.	Dependence on a fully known SCM/causal graph	[3],[5],[6], [7],[12],[13], [14],[15],[16], [17],[18],[20]
3.	Confounding and non-identifiability	[8],[19]

Feasibility and Actionability under Real-World Constraints

UNDER PEER REVIEW IN IJAR

This review begins with Diverse Counterfactual Explanations (DiCE) [2], a method that improves upon traditional counterfactual explanations by generating multiple diverse, feasible, and actionable counterfactuals instead of a single explanation, producing more realistic recommendations. However, it does not account for causal relationships or user preferences, which may limit its effectiveness.

Feasible and Actionable Counterfactual Explanations (FACE) [4] is another method that generates feasible and actionable paths rather than only suggesting a desired outcome. The authors argue that explanations should account for both the nature of the counterfactual and the user's ability to realistically achieve it. As future work, the authors suggest evaluating FACE on more real-world datasets.

Another method with the same limitation, but far more efficient, is Causal Algorithmic Recourse with Model-based Amortization (CARMA) [9], a neural-network-based approach that uses pretrained causal generative models to efficiently generate recourse recommendations with limited causal information.

Returning to the importance of the robustness and meaningfulness of generated counterfactuals, On the Adversarial Robustness of Causal Algorithmic Recourse [10] formulates the adversarially robust recourse problem and shows that existing minimal-cost recourse methods are fragile. The authors propose methods for generating causally robust recourse and a model regularizer that encourages classifiers to behave locally linearly and rely on actionable features. The limitation is that their approach trades off recourse cost against robustness, and future work is needed to extend robust recourse generation to broader forms of uncertainty and less constrained classifiers.

Another problem with counterfactuals is that they may become ineffective over time as the world around them keeps changing, as shown in Time Can Invalidate Algorithmic Recourse [11]; using a causal lens, the authors demonstrate this failure mode and propose an algorithm to address it.

Dependence on a Fully Known SCM/Causal Graph

This gap is partly addressed by algorithmic recourse [3], which treats recommended actions as causal interventions and seeks the minimal feasible intervention needed to obtain a favorable prediction - reflecting a broader shift in XAI toward telling individuals how a favorable outcome can be achieved, rather than only explaining a decision after the fact. Its main limitation is that it assumes access to an accurate causal model of realistic settings; suggested future work includes developing better causal world models and extending the approach to more complex settings.

UNDER PEER REVIEW IN IJAR

A related concept, causal constraints, refers to constraints embedded within a model to ensure that a counterfactual explanation is feasible and consistent with practical cause-and-effect relationships. This is explored in Preserving Causal Constraints in Counterfactual Explanations for Machine Learning Classifiers [15], which proposes two approaches: one uses Structural Causal Models (SCMs) to preserve causal constraints during counterfactual generation, while the other learns feasibility constraints from user feedback using a Variational Autoencoder (VAE) when causal knowledge is unavailable. The main limitation is that modelling feasibility remains challenging, especially when causal knowledge is incomplete.

A similar limitation is observed in Counterfactual Explanations as Interventions in Latent Space (CEILS) [5], a framework that incorporates causal relationships into counterfactual generation while remaining compatible with existing counterfactual methods. Rather than replacing current generators, CEILS improves their outputs by producing more feasible and actionable counterfactual explanations that respect causal dependencies. Future work includes handling partial or uncertain causal graphs and relaxing SCM assumptions, since this ideal scenario is uncommon in real-world models.

To address this limitation, two probabilistic approaches [6] generate interventions with a high probability of achieving recourse under limited causal knowledge. One models uncertainty in the structural equations using Gaussian Processes, while the other estimates intervention effects using similar subpopulations. Both approaches account for an imperfect causal graph, reflecting the non-ideal conditions under which counterfactuals must typically be computed. However, the proposed methods still rely on assumptions about the available causal information and uncertainty model, and future work aims to improve recourse under even more limited causal knowledge.

All of the above approaches and methods share one thing in common: they assume the availability of an underlying, perfect SCM and causal graph, which is generally not the case. This review now turns to an approach that attempts to address this limitation [7], where the authors propose the first framework that generates feasible causal recourse without relying on a predefined causal graph or SCM. Recognizing that causal counterfactuals cannot be reliably identified from observational data alone, they reformulate algorithmic recourse as structure-preserving interventions and approximate causal relationships using Independent Component Analysis (ICA). Their methods preserve causal dependencies while producing more feasible, interpretable, and low-cost recourse recommendations. The limitation here is that the framework still assumes causal sufficiency and does not fully address unobserved confounders. Future work aims to incorporate auxiliary information and extend the method to settings with hidden confounding.

To address this same limitation of not knowing the SCM, another proposed method uses the Structural Causal Marginal Problem (SCMP) [20], a framework that enables causal and

UNDER PEER REVIEW IN IJAR

counterfactual inference by combining information from multiple observational and experimental datasets instead of requiring a single complete dataset. The limitation here is that the framework depends on the availability of compatible partial datasets and assumes they can be consistently combined; nonetheless, it relaxes the assumption of an underlying SCM and causal graph, which is precisely the goal of this line of research. Future work focuses on extending SCMP to more complex causal systems and improving methods for integrating diverse causal information while maintaining reliable counterfactual inference.

Another approach, causal imputation [12], uses partial knowledge about a causal system to predict outcomes under different possible actions. It extends SCM based models to situations where the observed outcome can be represented as a counterfactual. The main limitation of this method is that it focuses mainly on causal prediction rather than directly addressing actionable counterfactual recourse; future work includes applying the framework to noisy and more complex settings. A related method, Advancing Counterfactual Inference through Nonlinear Quantile Regression [13], addresses a similar constraint: the difficulty of performing counterfactual inference without access to a predefined SCM. It proposes using observational data and learned causal structure by formulating counterfactual inference as a quantile regression problem. Its main limitation, again, is that it still largely depends on the learned causal structure, and future work can aim to improve reliability under incomplete causal knowledge.

Following the trend of relaxing the assumption of a known SCM, the rank-preservation approach [14] identifies counterfactual outcomes and introduces a learning method with theoretical guarantees for unbiased and efficient estimation, but this method becomes unstable in settings with sparse data; future work aims to improve the framework under such conditions.

A method that partly solves the problem of relying on a complete, predefined SCM is Conditional Intervention for Algorithmic Recourse (CIAR) [16]; it solves the problem only partly, as it still relies on some causal information, but it demonstrates improved effectiveness and efficiency compared with previous approaches.

Apart from identifying counterfactuals, it is also necessary to ensure that they are actually useful in real-world scenarios, that is, meaningful and robust. Traditional algorithmic recourse may change a model's prediction without actually improving the practical outcome. To address this, one proposed approach is Meaningful Algorithmic Recourse (MAR), which recommends actions that improve both the prediction and the underlying target, and Endogenous Algorithmic Recourse (EAR), which limits interventions to causes of the target [17]. The limitation here is the same as in many papers discussed above: the assumption of complete knowledge of the SCM and causal relationships.

Similarly, Improvement-Focused Causal Recourse (ICR) [18] works toward generating genuinely helpful counterfactuals and is one of the few studies that aims at real improvement rather than mere acceptance.

UNDER PEER REVIEW IN IJAR

Confounding and Non-Identifiability

Another limitation common in such cases is known as confounding: it occurs when an unobserved variable influences both the cause (intervention) and the outcome, creating a misleading association between them. As a result, the true causal effect of an intervention cannot be accurately estimated. Since hidden confounders are common in practical settings, these assumptions can lead to unreliable recourse recommendations. The authors [8] propose a framework that relaxes these assumptions by accounting for unobserved confounding and computing bounds on counterfactual effects rather than exact estimates. The main limitation here is that the approach still relies on knowledge of the causal graph and confounding structure, which may not always be available. Future work includes improving the computational efficiency of the algorithms and extending the framework to broader fairness and real-world applications.

Alongside this method, another approach learns an SCM using observational data alone, but with a major limitation: learning an SCM from observational data does not necessarily guarantee reliable counterfactual inference. This is demonstrated in Counterfactual (Non-)identifiability of Learned Structural Causal Models [19], which proves that identifiability holds only under certain simple conditions, such as monotonic, single-dimensional generation mechanisms, and proposes a worst-case error estimation method that quantifies the reliability of counterfactual predictions, allowing practitioners to assess how much a learned Deep Structural Causal Model (DSCM) can be trusted. The limitation primarily highlighted in this study is that the proposed approach does not resolve the fundamental problem of counterfactual non-identifiability for general multi-dimensional generation mechanisms; instead, it estimates the uncertainty associated with the predictions. It does not necessarily provide a new generative method but rather explores an existing limitation; future work aims to develop methods that improve the identifiability of learned SCMs, extend the analysis to more complex causal settings, and provide stronger guarantees for reliable counterfactual inference.

TABLE II. Taxonomy of Counterfactual Explanations and Algorithmic Recourse Methods.

S.NO	Method	Causal Assumption	Feasibility Type	Main Limitation
1.	Counterfactual Explanations without Opening the Black	No explicit causal model; relies on feature changes to	Counterfactual explanations based on changes to	May suggest changes that are not actionable or causally feasible in

UNDER PEER REVIEW IN IJAR

	Box [1]	explain individual predictions	input features	the real world
2.	DiCE [2]	Does not require an explicit causal model	Diverse counterfactuals subject to feature constraints	Does not explicitly account for causal relationships or user preferences
3.	Algorithmic Recourse [3]	Assumes access to an accurate causal model	Minimal feasible causal interventions	Depends on the availability and accuracy of the causal model
4.	FACE [4]	Uses the data distribution to identify feasible paths; no explicit causal model required	Data-manifold-based feasible and actionable paths	Feasible paths are not necessarily causally valid
5.	CEILS[5]	Incorporates causal relationships into counterfactual generation	Causally consistent, feasible and actionable counterfactuals	Handling partial causal graphs and relaxing SCM assumptions remains challenging
6.	Probabilistic Recourse [6]	Uses probabilistic causal models and allows imperfect causal knowledge	Recourse under causal uncertainty	Relies on assumptions about causal information and uncertainty
S.NO	**Method**	**Causal Assumption**	**Feasibility Type**	**Main Limitation**
7.	Causal Recourse without Prior Structural Knowledge [7]	Does not require a predefined causal graph or SCM; assumes causal sufficiency	Feasible, structure-preserving causal recourse	Does not fully address unobserved confounding
8.	Recourse under Confounding [8]	Allows partially or fully confounded settings	Bounds on counterfactual effects	Requires information about the causal graph and confounding structure
9.	CARMA [9]	Uses pretrained causal generative models	Scalable recourse recommendations	Depends on the available causal information and model setup
10.	Adversarially Robust Recourse [10]	Uses a causal model and an adversarial threat mode	Robustness to adversarial changes	Trades off recourse cost against robustness
11.	Temporal Recourse [11]	Considers changes in the environment over time	Recourse that accounts for temporal changes	Previously valid recourse can become ineffective as conditions change
12.	Causal Imputation [12]	Uses partial knowledge of the causal system	Counterfactual outcome prediction	Focuses on causal prediction rather than directly generating actionable recourse
13.	Nonlinear Quantile Regression [13]	Learns causal structure from observational data	Counterfactual inference using quantile regression	Reliability still depends on the learned causal structure

UNDER PEER REVIEW IN IJAR

14.	Rank-Preservation Approach [14]	Assumes rank preservation for counterfactual outcome estimation	Counterfactual outcome estimation	Can be unstable with sparse data
15.	Causal Constraints [15]	Uses an SCM or learned feasibility constraints	Constraint-based counterfactual feasibility	Feasibility modeling remains challenging when causal knowledge is incomplete
16.	Conditional Intervention [16]	Uses some causal information for conditional interventions	Intervention-based algorithmic recourse	Does not eliminate dependence on causal information
17.	Meaningful and Robust Recourse [17]	Uses a causal model to assess meaningfulness and robustness	Meaningful and robust recourse	Assumes knowledge of the SCM and causal relationships
18.	Improvement-Focused Causal Recourse (ICR) [18]	Uses causal recourse to target practical improvement	Improvement-focused recourse	Dependence on known SCM
S.NO	**Method**	**Causal Assumption**	**Feasibility Type**	**Main Limitations**
19.	Counterfactual Non-Identifiability [19]	Studies identifiability in learned SCMs	Identifiability and reliability assessment	General multidimensional counterfactuals may remain non-identifiable
20.	Structural Causal Marginal Problem (SCMP) [20]	Combines compatible partial observational and experimental information	Causal inference with incomplete datasets	Depends on the availability and compatibility of partial datasets

SUMMARY

This review traces a clear progression in the development of counterfactual explanations and algorithmic recourse. Earlier papers introduced counterfactual explanations as a way of providing actionable recourse rather than plain explanations, while later methods such as DiCE [2] and FACE [4] improved on these by making recommendations more actionable, feasible, and realistic - moving the field toward the causality of counterfactuals.

UNDER PEER REVIEW IN IJAR

As the review progresses, researchers are seen incorporating SCMs, causal graphs, and causal relationships to align counterfactuals with real-world settings [15], [5], but then confronting the reality that a complete underlying SCM or full causal information is usually unavailable, prompting studies to relax these assumptions while still generating actionable and feasible counterfactuals [5] ,[6], [7], [12],[13],[14],[20] Another realistic problem is that of hidden confounders, which many studies aim to solve [8], and most recently the field has shifted focus toward the scalability, adversarial robustness, and temporal stability of recourse recommendations [9],[10],[11],[18]. Across all of this work, a common thread emerges: the field has progressively relaxed the idealized assumption of a fully known and correct causal model, replacing it with methods that operate under partial knowledge and bounded uncertainty.

However, no method reviewed here fully resolves feasibility, confounding, and incomplete causal knowledge simultaneously, as each new method or contribution tends to relax one assumption while retaining the others - suggesting that an open direction for future research is addressing these three gaps simultaneously.

DISCUSSION AND FUTURE DIRECTION

The study focuses on three major research gaps and demonstrates different ways in which these can be resolved. For the first gap, "Feasibility and Actionability," rather than examining factors such as validity over time, practical implementation, and generation in isolation, future researchers could observe how these factors act in accordance with one another and how one affects the other, in order to arrive at a complete method that resolves all of these limitations jointly. For "Dependence on a Fully Known SCM," rather than completely relaxing the assumption at once, one could start with partial information about the SCM and progressively improve the method to eventually eliminate the dependence entirely. For "Confounding and Non-Identifiability," future work could address these two problems jointly rather than in isolation, since confounding and non-identifiability frequently co-occur in real-world causal systems and are unlikely to be fully resolved when treated as independent limitations.

The study also observes a recurring trend across the different papers: each paper focuses on relaxing only a single assumption, that is, it addresses just one limitation at a time. In the future, researchers should aim to incorporate all three gaps, and potentially the different methods proposed to address them, into an integrated framework capable of solving all three problems simultaneously. This task may prove quite difficult, as these gaps are individually challenging to

UNDER PEER REVIEW IN IJAR

overcome, but it could become a viable future direction once further advancements are made as the field progresses.

CONCLUSION

In conclusion, this paper analyzed the progress of counterfactuals in recent years and identified research gaps and different problems by reviewing 20 carefully selected research papers. Three major gaps were identified and explained: the first being "Dependence on the SCM," as a fully defined SCM is usually not available in practical, real-world scenarios; second, "Confounding and Non-Identifiability," which represents the problem that many hidden variables may contribute to the recourse; and finally, "Feasibility and Actionability under Real-World Constraints," that is, the recourse must be possible for the user to carry out in practical settings, as counterfactuals have many practical and important applications in industries such as healthcare, banking, finance, and other important sectors.

This review, in synthesizing the progress and limitations of twenty studies spanning 2018 to 2026, aims to serve as a consolidated reference point for researchers entering this field, as well as a starting direction for future work seeking to address feasibility, SCM-dependence, and confounding. The study can also provide foundational knowledge to support improvements in the field. It is also observed that, across the papers reviewed, there is a trend where each new method tends to relax one assumption but still largely depends on the others; therefore, for future work, researchers should aim to integrate different methods, or create new methods, that relax the assumptions of all three gaps discussed above.

AI USE DISCLOSURE

ChatGPT (OpenAI) and Claude (Anthropic) were used to assist with proofreading, citation formatting, and basic document formatting. The author reviewed the suggested changes and remains responsible for the accuracy, interpretation, and final content of the pape

REFERENCES

[1] S. Wachter, B. Mittelstadt, and C. Russell, "Counterfactual Explanations without Opening the Black Box: Automated Decisions and the GDPR," Harvard Journal of Law & Technology, vol. 31, 2018.

UNDER PEER REVIEW IN IJAR

[2] R. K. Mothilal, A. Sharma, and C. Tan, "Explaining Machine Learning Classifiers through Diverse Counterfactual Explanations," in Proc. ACM Conf. Fairness, Accountability, and Transparency (FAT*), 2020.

[3] A.-H. Karimi, B. Schölkopf, and I. Valera, "Algorithmic Recourse: From Counterfactual Explanations to Interventions," in Proc. ACM Conf. Fairness, Accountability, and Transparency (FAccT), 2021, pp. 353–362.

[4] R. Poyiadzi, K. Sokol, R. Santos-Rodriguez, T. De Bie, and P. Flach, "FACE: Feasible and Actionable Counterfactual Explanations," in Proc. AAAI/ACM Conf. AI, Ethics, and Society (AIES), 2020, pp. 344–350.

[5] R. Crupi, A. Castelnovo, D. Regoli, and B. San Miguel Gonzalez, "Counterfactual explanations as interventions in latent space," *Data Mining and Knowledge Discovery*, vol. 38, no. 5, pp. 2733–2769, 2024, doi: 10.1007/s10618-022-00889-2.

[6] A.-H. Karimi, J. von Kügelgen, B. Schölkopf, and I. Valera, "Algorithmic Recourse under Imperfect Causal Knowledge: A Probabilistic Approach," in Advances in Neural Information Processing Systems (NeurIPS), vol. 33, 2020, pp. 265–277.

[7] H. Wang, H. Zou, X. Zhou, S. Wang, W. Yang, and P. Cui, "Learning feasible causal algorithmic recourse: A prior structural knowledge free approach," in *Proc. ACM Web Conf. (WWW '25)*, 2025, doi: 10.1145/3696410.3714859.

[8] J. von Kügelgen, N. Agarwal, J. Zeitler, A. Mastouri, and B. Schölkopf, "Algorithmic Recourse in Partially and Fully Confounded Settings Through Bounding Counterfactual Effects," arXiv preprint arXiv:2106.11849, 2021.

[9] A. Majumdar and I. Valera, "CARMA: A Practical Framework to Generate Recommendations for Causal Algorithmic Recourse at Scale," in Proc. ACM Conf. Fairness, Accountability, and Transparency (FAccT), 2024, pp. 1745–1762.

[10] R. Dominguez-Olmedo, A.-H. Karimi, and B. Schölkopf, "On the Adversarial Robustness of Causal Algorithmic Recourse," in Proc. Int. Conf. Machine Learning (ICML), 2022, pp. 5324–5342.

[11] G. De Toni, S. Teso, B. Lepri, and A. Passerini, "Time Can Invalidate Algorithmic Recourse," arXiv preprint arXiv:2410.08007, 2024.

[12] Á. Ribot, C. Squires, and C. Uhler, "Causal Imputation for Counterfactual SCMs: Bridging Graphs and Latent Factor Models," in Proc. Conf. Causal Learning and Reasoning (CLeaR), 2024, pp. 1141–1175.

UNDER PEER REVIEW IN IJAR

[13] S. Xie, B. Huang, B. Gu, T. Liu, and K. Zhang, "Advancing Counterfactual Inference through Nonlinear Quantile Regression," arXiv preprint arXiv:2306.05751, 2023.

[14] P. Wu, H. Li, C. Zheng, Y. Zeng, J. Chen, Y. Liu, R. Guo, and K. Zhang, "Learning Counterfactual Outcomes Under Rank Preservation," arXiv preprint arXiv:2502.06398, 2025.

[15] D. Mahajan, C. Tan, and A. Sharma, "Preserving Causal Constraints in Counterfactual Explanations for Machine Learning Classifiers," arXiv preprint arXiv:1912.03277, 2019.

[16] Chiung-Yu Chang, Keng-Te Liao, Mi-Yen Yeh and Shou-De Lin, "Fast Conditional Intervention in Algorithmic Recourse with Reinforcement Learning," submitted to the International Conference on Learning Representations (ICLR), 2024. [Online]. Available: https://openreview.net/forum?id=oVVLBxVmbZ.

[17] G. König, T. Freiesleben, and M. Grosse-Wentrup, "A Causal Perspective on Meaningful and Robust Algorithmic Recourse," arXiv preprint arXiv:2107.07853, 2021.

[18] G. König, T. Freiesleben, and M. Grosse-Wentrup, "Improvement-Focused Causal Recourse (ICR)," in Proc. AAAI Conf. Artificial Intelligence, 2023.

[19] A. Nasr-Esfahany and E. Kiciman, "Counterfactual (Non-)identifiability of Learned Structural Causal Models," arXiv preprint arXiv:2301.09031, 2023.

[20] L. Gresele, J. von Kügelgen, J. M. Kübler, E. Kirschbaum, B. Schölkopf, and D. Janzing, "Causal Inference Through the Structural Causal Marginal Problem," in Proc. Int. Conf. Machine Learning (ICML), 2022.

Corresponding Author:-
Address:-

UNDER PEER REVIEW IN IJAR

UNDER PEER REVIEW IN IJAR