Thank you for the thoughtful, detailed, and overall very positive feedback. Reviewers agree that the paper is *well written* (**R1**, **R2**, **R3**), *well motivated* (**R1**), *clear* (**R2**), and *novel* (**R1**, **R2**). They consider the topic *important* (**R2**) and our methodology *sound* (**R2**, **R3**), *technically as well as practically interesting*, and *of independent interest* (**R1**). Reviewers find *the experimental evaluation extensive* (**R1**) and *well described* (**R2**), *our empirical evidence clear*, and *claims well supported and correct* (**R2**). We address the main concerns below and will incorporate all reviewers' feedback.

R1 (+R3): *Recourse is a counterfactual query, but the proposed solutions are based on counterfactual as well as interventional queries. Which framing is more appropriate?* We mostly agree that recourse is predominantly a counterfactual query. This is certainly the case in static environments where any changes result from an individual's actions. In dynamic settings, however, the unobserved SCM background variables U_r can also capture environmental factors which may change, e.g., between consecutive loan applications, and should thus not be kept fixed as in the counterfactual approach. Depending on the setting, recourse may thus constitute a mixture of interventional and counterfactual queries. Independently of this fundamental question, our interventional notion of recourse is practically motivated by a reluctance to make (parametric) assumptions about the structural equations which are both untestable but also necessary for counterfactual reasoning (see also Example 1; l.118–120; 181–186). In such cases, the subpopulation-based approach can provide a useful alternative. In analogy, in epidemiology one ideally wants to estimate the individualised (i.e., counterfactual) treatment effect, but (without additional assumptions) has to resort to the (conditional) average (i.e., interventional) effect instead. Finally, we note that the proposed approaches can be combined depending on the available domain knowledge since each causal relation is treated separately.

R1, R3: *Lack of comparison to existing baselines and qualitative assessment of the type of recourse that the model learns.* We compare with an oracle (known SCM) corresponding to [7], as well as with point-based linear and kernel-ridge regression baselines. A comparison with approaches assuming independent features raises issues such as whether descendant variables are allowed to change or whether keeping their value constant should incur a cost (and if so how much), and are thus neither natural nor straight-forward. We therefore restrict ourselves to a comparison with *causal* recourse approaches. For a qualitative assessment of the recommended actions, we refer to Table 5 in Appendix E.

R1: *Does the learned CVAE respect the causal graph? The procedure of training CVAEs for interventional recourse should be clarified.* Yes, the interventional approaches respect the causal graph: we do not train a single CVAE model, but—motivated by the identifiability result (11) in Proposition 4—instead learn a separate CVAE/GP for each parent-child conditional $p(X_r|\mathbf{X}_{\text{pa}(r)})$, see also l.193–201. We apologise for failing to convey this important point.

R1: *Can the authors clarify the enumeration of feasible actions? There are causal dependencies in the graph which can determine allowable feasible sets.* We agree that feasibility is an essential aspect of recourse. However, our understanding is that causal dependencies in the graph *do not* determine allowable feasible sets. E.g., consider a mutable but non-actionable variable `BMI` (body mass index), a descendant of an actionable variable `weight` and an immutable variable `gender`. The mutability and actionability constraints above are determined by context-/individual-specific constraints in the application domain (sumarised in $\mathbb{A}^{\text{F}}$), rather than by the mathematical framework of SCMs.

R1: *The authors do fail to cite an important critique of recourse work [VA20][1] and contextualize their contributions in terms of issues raised in this work.* Thank you for the pointer. Indeed, concerns as the ones raised in [VA20] were a primary motivation for our paper. Specifically, [VA20] explore a number of challenges facing algorithmic recourse, and argue (alongside [BSR20][2]) that we need to look beyond independently manipulable features; [7] then phrases this as a *causal* problem within the SCM framework. We apologize for failing to cite [VA20, BSR20] and will correct this.

R2: *It is not clear whether "action" and "intervention" refer to slightly different concepts or not.* Thank you for pointing this out; we will revise the wording to clarify that we consider actions and interventions as interchangeable.

R3: *How do computational complexity of probabilistic and point-based approaches compare?* Their main difference is the expectation in the constraint whose approximation scales linearly in the number of Monte Carlo samples.

R3: *Some parts of the results ($\mathcal{M}_{\text{KR}}$ on nonlinear ANM, $\text{CATE}_{\text{CVAE}}$ vs $\mathcal{M}_{\text{CVAE}}$, non-additive SCM) are not convincing. The paper should include a scenario where subpopulation-based approaches are preferable.* The strong performance of $\mathcal{M}_{\text{KR}}$ on the nonlinear ANM is expected as its assumptions are satisfied. Moreover, $\text{CATE}_{\text{CVAE}}$ indeed outperforms $\mathcal{M}_{\text{CVAE}}$ on the same. We agree regarding the non-additive SCM (where subpopulation-based approaches should be preferred) and (i) found a more convincing example by also making $X_3|X_1, X_2$ multimodal. In general, misspecified models tend to intervene on leaf nodes (due to failure to capture causal relations) and thereby can still achieve recourse, though at higher cost. We also tried (ii) making leaf nodes non-actionable (such as `BMI` from the example above) and (iii) reducing the cost of acting on root nodes, both of which also make the advantageous cost-validity trade-off of our methods more pronounced. We will include these additional findings in the revised version.

Overall, we hope that the above will convince all reviewers that *"this is a significant contribution [in] the area of actionable recourse"* (**R1**) and that it *"is definitely relevant to the NeurIPS audience"* (**R1**, **R2**).

---

[1] [VA20] Venkatasubramanian, Alfano. "The philosophical basis of algorithmic recourse." FAccT 2020.

[2] [BSR20] Barocas, Selbst, Raghavan. "The hidden assumptions behind counterfactual explanations and principal reasons." FAccT 2020.