

Kernel Distributionally Robust Recourse Action

In the current era, machine learning models are increasingly deployed across nearly all high-stakes domains, for instance, in banking for loan approvals, in healthcare for treatment recommendations or disease risk assessment, in recruitment for candidate screening, and in resource allocation for critical decision-making. Within this context, recourse in machine learning refers to the set of actionable modifications an individual can make in order to obtain a favorable outcome[1], given that their current profile results in an unfavorable decision. For example, consider the case of a loan application: if a bank rejects an applicant due to their present salary, investments, and property status, a recourse action might suggest increasing the monthly salary by a certain percentage or improving investment levels so that the application is reconsidered and ultimately approved. By providing personalized and actionable recommendations, recourse enhances not only individual agency but also trust in machine learning systems.

Most existing methods for generating recourse rely on the assumption of a fixed predictive model. In practice, however, this assumption rarely holds. Data distributions are inherently dynamic, evolving due to population shifts, domain transfer, or temporal drift. As a result, recourses generated under fixed-model assumptions often become invalid or unreliable. Robust counterparts based on interval uncertainty and Wasserstein ambiguity sets[2] have been proposed to address this issue. While effective in linear or quadratic settings, these methods fail to capture nonlinear interactions and often lead to conservative recommendations when applied to complex, high-dimensional data.

To bridge this gap, we introduce Kernel Distributionally Robust Recourse Action (KDRRA), a novel framework that leverages Maximum Mean Discrepancy (MMD) to model nonlinear distributional shifts within a reproducing kernel Hilbert space (RKHS)[3]. Our approach combines kernel methods with duality theory, yielding a tractable reformulation in the form of a second-order convex program. This formulation makes it possible to compute recourses that remain valid even under complex, nonlinear shifts.

Furthermore, we extend our framework by proposing a Bayesian Kernel DRO formulation. Unlike standard kernel DRO, which relies heavily on the empirical distribution and may require overly large ambiguity radii to ensure robustness, the Bayesian extension integrates posterior beliefs about the data-generating process. This reduces unnecessary conservatism, ensuring that the recourse recommendations remain both reliable and practical.

We conduct extensive numerical experiments on three real-world datasets (German, SBA, Student), benchmarking our method against state-of-the-art recourse approaches. The results consistently demonstrate that KDRRA achieves higher validity of recourses under distributional shifts, while maintaining competitive or lower costs compared to existing robust methods. The Bayesian extension further improves stability when training data are limited or noisy. Together, these contributions highlight the potential of KDRRA as a principled and scalable framework for reliable recourse in dynamic decision-making environments.

Keywords: Actionable Recourse, Distributionally Robust Optimization (DRO), Kernel Methods, RKHS, Maximum Mean Discrepancy (MMD), Bayesian DRO.

References:

- 1 Ustun, B., Spangher, A. and Liu, Y., 2019, January. Actionable recourse in linear classification. In Proceedings of the conference on fairness, accountability, and transparency (pp. 10-19).
- 2 Nguyen, D., Bui, N. and Nguyen, V.A., 2023. Distributionally robust recourse action. arXiv preprint arXiv:2302.11211.
- 3 Yang, S.B. and Li, Z., 2022. Kernel distributionally robust chance-constrained process optimization. Computers & Chemical Engineering, 165, p.107953.