From Individual to Multi-Agent Algorithmic Recourse: Minimizing the Welfare Gap via Capacitated Bipartite Matching

Anonymous Author(s)
Affiliation
Address
`email`

Abstract

Decision makers are increasingly relying on machine learning in sensitive situations. In such settings, algorithmic recourse aims to provide individuals with actionable and minimally costly steps to reverse unfavorable AI-driven decisions. While existing research predominantly focuses on single-individual (i.e., seeker) and single-model (i.e., provider) scenarios, real-world applications often involve multiple interacting stakeholders. Optimizing outcomes for seekers under an individual welfare approach overlooks the inherently multi-agent nature of real-world systems, where individuals interact and compete for limited resources. To address this, we introduce a novel framework for multi-agent algorithmic recourse that accounts for multiple recourse seekers and recourse providers. We model this many-to-many interaction as a capacitated weighted bipartite matching problem, where matches are guided by both recourse cost and provider capacity. We propose a three-layer optimization framework: (1) basic capacitated matching, (2) optimal capacity redistribution to minimize the welfare gap, and (3) cost-aware optimization balancing welfare maximization with capacity adjustment costs. Experimental validation demonstrates that our framework enables near-optimal welfare with minimal system modifications.

1 Introduction

AI decision-making systems rapidly apply predictive models to support individuals in various contexts, e.g., loan approvals, medical treatments, or bail decisions [9]. Driven by AI policy regulations and the idea of a "right to explanation," algorithmic recourse is an emerging field that aims to provide individuals affected by negative, high-stakes algorithmic decisions with recommendations on how to reverse those outcomes [2, 8]. Therefore, algorithmic recourse refers to the systematic process of reversing unfavorable decisions made by algorithms across various counterfactual scenarios [10]. Existing studies on algorithmic recourse predominantly address how the individual would need to change their attributes to achieve the desired outcome [4]. Such settings generally assume a single individual impacted by a single decision-making model as shown in Figure 1a.

In real-world scenarios, however, AI decision-making systems (i.e., *providers*) often interact with multiple individuals whose actions can influence outcomes and, consequently, recourse recommendations for others. Furthermore, individuals seeking recourse (i.e., *seekers*) may engage with multiple providers (Figure 1c) to choose the most suitable among given recommendations. While these studies have extended the literature to settings with multiple recourse seekers [6, 1], they continue to assume a single provider. There remains a gap regarding algorithmic recourse in situations involving multiple recourse providers each potentially impacting outcomes with their own decision models. Existing approaches typically overlook how providing recourse recommendations simultaneously to multiple

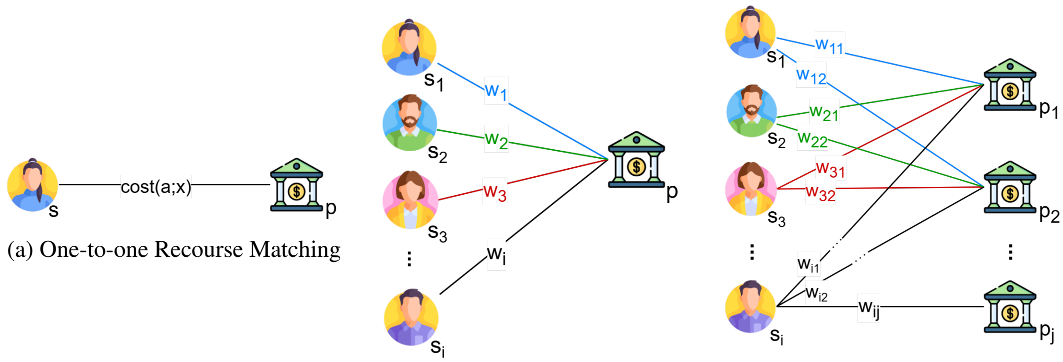

(a) One-to-one Recourse Matching

(b) Many-to-one Setting

(c) Many-to-many Recourse Matching

Figure 1: Various algorithmic recourse setups: (a) the one-to-one setting, where an individual s seeks recourse recommendations from a provider p with minimal cost required to reverse the output; (b) The many-to-one setting, where multiple individuals are seeking recourse from a single recourse provider. (c) our proposed many-to-many framework generalizes prior settings by simultaneously optimizing for multiple recourse seekers and providers.

recourse seekers can benefit society and overall recourse actionability through interactions among individuals.

We propose a framework that includes multiple recourse seekers and multiple recourse providers. We formalize this interaction as a *capacitated weighted bipartite matching problem* and determine optimal recourse outcomes using a linear-programming approach, thereby maximizing social welfare under capacity constraints. Further, we identify a welfare gap between the socially optimal solution, computed by a central planner, and the unrealistic individually optimal outcome, where each seeker acts without consideration of providers' capacity constraints. To minimize this gap, we introduce a second optimization layer that finds the best distribution for a total fixed capacity over providers. Finally, we add the third optimization layer that minimizes the welfare gap while penalizing deviations from the initial capacity values.

Our framework shows how uncoordinated individual recourse decisions can lead to collective inefficiencies. Also, our capacity redistribution approaches align with collective goals and strategies, coordination challenges, and the design of mechanisms for socially beneficial outcomes.

2 Many-to-Many Recourse Optimization

We formalize the matching problem as a bipartite graph with set of seekers $\mathcal{S}$: $\{s_i \mid s_i \in \mathcal{S}, \forall\, i \in [n]\}$, each characterized by a feature vector x_i and set of providers $\mathcal{P}$: $\{p_j \mid p_j \in \mathcal{P}, \forall\, j \in [m]\}$, each equipped with a classifier (w.l.o.g. binary model) h_j to accept or reject the seekers and a matching capacity k_j. All seekers are initially rejected by all providers. Furthermore, it is assumed that a *central planner* will coordinate matches between seekers and providers (i.e., eq. (1)) and potentially redistribute existing capacity among providers (i.e., eq. (2) and eq. (3)).

Recourse cost computation As defined by Ustun et al. [7], given provider p_j's decision model h_j and an input feature vector x_i corresponding to the characteristics of seeker s_i, such that $h_j(x_i) = -1$ (assumed binary w.l.o.g.), the recourse cost for seeker s_i to achieve approval from provider p_j is the solution to the optimization problem as below:

$$c_{ij} = \min_{a \in A(x_i)} \text{cost}(a; x_i) \quad \text{s.t.} \quad h_j(x_i + a) = +1 \quad \forall\, i, j$$

where a is an action vector representing feasible changes to the features of x_i, and $A(x_i)$ is the set of allowed actions based on domain constraints (e.g., mutability and bounds on feature changes). Our proposed framework is agnostic to the choices of recourse method and providers' model, operating only on the minimum cost of change required for each pair of seeker s_i and provider p_j.

Bipartite Graph Construction. We construct a *weighted bipartite graph* $\mathcal{G} = (\mathcal{V}, \mathcal{W})$, where nodes $\mathcal{V} = \mathcal{S} \cup \mathcal{P}$ and $\mathcal{W} := \{w_{ij} \mid w_{ij} = e^{-\gamma \cdot c_{ij}}, \forall\, i, j\}$. where $\gamma > 0$ is a scaling parameter controlling the sensitivity of the transformation. This exponential transformation converts costs into edge weights, enabling algorithms such as the maximum-weight bipartite matching [5] to prioritize low-cost (i.e., efficient) recourse assignments while maximizing overall match coverage.

Optimization Model Next, *capacitated weighted bipartite matching problem.* If we denote the maximum weight for seeker i accordingly as $w_i^* = \max_j(w_{ij})$, we can then measure the ideal scenario in which each seeker attains its optimal outcome as:

$$\text{Individual Welfare} := \sum_{i=1}^{n} w_i^*$$

However, this ideal scenario assumes that providers have unbounded capacity (w.l.o.g., at least the number of seekers for each provider), meaning that they can freely provide the resources, which is not realistic. In practice, each provider has a limited capacity k_j, meaning they can serve only a finite number of seekers. Taking a systems-level view and aiming to minimize the overall cost of recourse across all seekers and providers,[1] To obtain the optimal matching, we formulate a mixed-integer linear program (MILP) and solve it with the Gurobi Optimizer [3][2] as follows:

$$\begin{aligned} \text{Social Welfare} := \max_{z_{ij}} \quad & \sum_{i=1}^{n}\sum_{j=1}^{m} w_{ij}\, z_{ij} \\ \text{s.t.} \quad & \underbrace{\sum_{j=1}^{m} z_{ij} \le 1 \quad \forall i,}_{\text{Matching Constraint}} \quad \underbrace{\sum_{i=1}^{n} z_{ij} \le k_j \quad \forall j,}_{\text{Capacity Constraint}} \quad \underbrace{z_{ij} \in \{0,1\} \quad \forall i, j}_{\text{Edge Constraint}} \end{aligned} \tag{1}$$

where z_{ij} are binary decision variables that indicate whether seeker i is assigned to provider j. The capacity constraint may result in some seekers matching to a more costly match (lower weight w_{ij}). This discrepancy is quantified by the gap between the ideal individual welfare and the realized social welfare:

$$\text{Welfare Gap} := \left(\sum_{i=1}^{n} w_i^*\right) - \left(\sum_{i=1}^{n}\sum_{j=1}^{m} w_{ij}\, z_{ij}\right).$$

This gap highlights a critical design challenge: given a fixed total amount of provider capacity, how should these limited resources be distributed across providers to minimize the welfare gap? A naive uniform distribution of provider capacities may lead to significant welfare losses. In contrast, allocating more capacity to providers associated with lower recourse costs, can substantially reduce the welfare gap, even when the total capacity remains fixed. In the next section, we present an approach for optimizing capacity distribution to minimize this gap.

3 Minimize Welfare Gap

Under a fixed total provider capacity $K = \sum_{j=1}^{m} k_j$, the welfare gap can vary depending on how capacity is distributed among providers. In fact, for any given K, there is an optimal allocation of provider capacities k_j that minimizes this welfare gap. This observation leads to a new optimization problem involving two sets of decision variables namely, the integer variables $k_j \;\forall\, j$, representing provider capacities in the optimal solution, and the matching variables z_{ij}, as previously defined, indicating the best matching under the system settings.

$$\max_{z_{ij},\, k_j} \quad \sum_{i=1}^{n}\sum_{j=1}^{m} w_{ij}\, z_{ij} \quad \text{s.t.} \quad \sum_{j=1}^{m} k_j = K \quad \forall\, j \quad \text{Total Capacity Constraint} \tag{2}$$

[1]This assumes that providers do not have ulterior preferences affecting the matching process.

[2]Although the presence of binary decision variables renders the problem NP-hard, Gurobi's branch-and-bound engine—augmented with presolve, cutting-plane generation, and heuristic warm-starts–guarantees global optimality.

The matching, capacity, and edge constraints remain the same as before in Equation (1), with the additional constraint on the total capacity.

We propose a systematic method that assigns capacities based on each seeker's top-ranked matching weight, previously defined in Section 2 as w_i^*. Since at most K seekers can receive recourse, excluding the top K highest-weight edges directly increases the achievable welfare. Thus, the welfare of any feasible solution is bounded above by $\sum_{i \in \mathcal{S}_K} w_i^*$, where $\mathcal{S}_K$ denotes the set of seekers corresponding to the K highest-ranked edges. These insights lead to Algorithm 1, which offers a more direct and efficient alternative to MILP.

Algorithm 1 Optimal Capacity Distribution

Input: seekers $\mathcal{S}$, providers $\mathcal{P}$, weights w_{ij}, total capacity K
Output: provider capacities $k = (k_1, \dots, k_{|\mathcal{P}|})$
Initialize empty list $\mathcal{L}$
for each seeker $i \in \mathcal{S}$ **do**
    $w_i^* = max_j(w_{ij})$
    $j_i^* \leftarrow \arg\max_{j \in \mathcal{P}} \; w_i^*$ {best provider for seeker i}
    Append triple $(i,\ j_i^*,\ w_i^*)$ to $\mathcal{L}$
end for
Sort $\mathcal{L}$ in descending order of weight
Select the first K triples of $\mathcal{L}$ {top-K matches}
Initialize $k_j \leftarrow 0$ for all $j \in \mathcal{P}$
for each selected triple $(i,\ j,\ w)$ **do**
    $k_j \leftarrow k_j + 1$
end for
return capacity vector k

Capacity Redistribution with Penalizing Modifications In practice, recourse methods often operate within established configurations determined by existing organizational structures, resource availability, and operational constraints. Transitioning from the current provider capacity configuration to an optimal setup typically involves real-world adjustment costs. Let $\hat{k}_j$ denote the initial capacity of provider j, and $\tilde{k}_j$ represent the target capacity after configuration. The change in capacity $\Delta k_j = \tilde{k}_j - \hat{k}_j$ can penalize large changes. Integrating this penalty into our optimization leads to a multi-objective problem, balancing social welfare maximization with minimization of capacity adjustment penalties. The modified objective function is:

$$\text{Welfare} = \max_{z_{ij},\, k_j} \left(\sum_{i=1}^{n} \sum_{j=1}^{m} w_{ij} z_{ij} - \sum_{j=1}^{m} \beta_j |\Delta k_j| \right) \tag{3}$$

subject to all the same constraints previously defined in Equation (2), and $\beta_j \geq 0$ controls the penalty sensitivity for each of the providers accordingly. This enhanced formulation constitutes a MILP problem, solvable by recent versions of solvers like Gurobi [3]. Ultimately, this optimization simultaneously identifies the z_{ij} and k_j, clearly indicating how capacities should be adjusted from their initial configuration. The Equation (3) generalizes the previous two through the penalty parameter β_j. When $\beta_j = 0$, the formulation reduces to Equation (2), and when $\beta_j \to \infty$, it enforces fixed capacities(Equation (1)). Thus, β_j enables interpolation between these two extremes.

4 Conclusion and Future Work

We have introduced a many-to-many view of algorithmic recourse in which multiple seekers obtain recommendations from multiple decision-making models whose resources are limited. Further, we quantified the welfare gap between the socially optimal solution, computed by a central planner, and the individually optimal outcome, where each seeker acts in isolation and selects the provider offering the lowest recourse cost, without coordination. In future versions of this work, we will present empirical evaluations to assess how well the proposed framework reduces welfare losses across synthetic and real datasets. These results will help quantify the welfare achievable with minimal capacity adjustments and demonstrate the framework's practical relevance.

References

[1] Joao Fonseca, Andrew Bell, Carlo Abrate, Francesco Bonchi, and Julia Stoyanovich. Setting the Right Expectations: Algorithmic Recourse Over Time. In *Equity and Access in Algorithms, Mechanisms, and Optimization*, pages 1–11, October 2023. doi: 10.1145/3617694.3623251. URL `http://arxiv.org/abs/2309.06969`. arXiv:2309.06969 [cs].

[2] GDPR. Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation) (Text with EEA relevance), April 2016. URL `http://data.europa.eu/eli/reg/2016/679/oj/eng`. Legislative Body: EP, CONSIL.

[3] Gurobi Optimization, LLC. Gurobi Optimizer Reference Manual, 2024. URL `https://www.gurobi.com`.

[4] Amir-Hossein Karimi, Gilles Barthe, Bernhard Schölkopf, and Isabel Valera. A Survey of Algorithmic Recourse: Contrastive Explanations and Consequential Recommendations. *ACM Comput. Surv.*, 55(5):95:1–95:29, December 2022. ISSN 0360-0300. doi: 10.1145/3527848. URL `https://dl.acm.org/doi/10.1145/3527848`.

[5] H. W. Kuhn. The Hungarian method for the assignment problem. *Naval Research Logistics Quarterly*, 2(1-2):83–97, 1955. ISSN 1931-9193. doi: 10.1002/nav.3800020109. URL `https://onlinelibrary.wiley.com/doi/abs/10.1002\discretionary{}/nav.3800020109`.

[6] Andrew O'Brien and Edward Kim. Toward Multi-Agent Algorithmic Recourse: Challenges From a Game-Theoretic Perspective. *Proceedings of the International Florida Artificial Intelligence Research Society Conference*, 35, May 2022. ISSN 2334-0762. doi: 10.32473/flairs.v35i.130614. URL `https://researchdiscovery.drexel.edu/esploro/outputs/journalArticle/Toward-Multi-Agent-Algorithmic-Recourse-Challenges-From/991019173643504721`.

[7] Berk Ustun, Alexander Spangher, and Yang Liu. Actionable recourse in linear classification. In *Proceedings of the conference on fairness, accountability, and transparency*, pages 10–19, 2019.

[8] Sahil Verma, Varich Boonsanong, Minh Hoang, Keegan Hines, John Dickerson, and Chirag Shah. Counterfactual Explanations and Algorithmic Recourses for Machine Learning: A Review. *ACM Comput. Surv.*, 56(12):312:1–312:42, October 2024. ISSN 0360-0300. doi: 10.1145/3677119. URL `https://dl.acm.org/doi/10.1145/3677119`.

[9] Paul Voigt and Axel Von Dem Bussche. *The EU General Data Protection Regulation (GDPR)*. Springer International Publishing, Cham, 2017. ISBN 978-3-319-57958-0 978-3-319-57959-7. doi: 10.1007/978-3-319-57959-7. URL `http://link.springer.com/10.1007/978-3-319-57959-7`.

[10] S. Wachter, B. Mittelstadt, and C. Russell. Counterfactual explanations without opening the black box: automated decisions and the GDPR. *Harvard Journal of Law and Technology*, 31(2), 2018. ISSN 0897-3393. URL `https://ora.ox.ac.uk/objects/uuid:86dfcdac-10b5-4314-bbd1-08e6e78b9094`. Publisher: Harvard Law School.