QUOTA CONSTRAINTS FOR DIVERSITY INTERVENTIONS IN SUBSET SELECTION

Anonymous authors
Paper under double-blind review

ABSTRACT

The combinatorial optimization problem of subset selection is often modeled as maximizing a set function that captures inter-element dependencies under some capacity/matroid constraints. In this paper, we examine this problem under "quota constraints" where the selected subset must meet some minimum group-wise quotas. We provide algorithms with guarantees for two popular scenarios extended to the quota-constrained setting and make an empirical case for their applicability to fair subset selection.

1 INTRODUCTION

The subset selection problem is found in various domains: hiring, choosing a cohort of candidates from a pool of people, recommending a small set of news articles, summarizing data, building a financial portfolio, etc. where one tries to choose an often capacity-constrained set $S \subset V$ (the universal set) so as to maximize a utility function $u(S)$. Very often, this utility is also strongly correlated with how *representative* the chosen subset is, of the data at large. As is usually done (Celis et al., 2016; Hesabi et al., 2015), we will break down the *diversity* of a set S into two notions.

Diversity that is captured in the utility function u. e.g. when the objective rewards geometric diversity over some feature space, or when u is submodular (Krause & Golovin, 2014; Lin & Bilmes, 2011; Chen et al., 2018; Feige & Izsak, 2013), i.e. it captures the diminished marginal gain in utility from adding a member to S when another similar member is already present in S.
Diversity that is enforced as an intervention. This includes constraints imposed on the optimization problem in order to overcome biases in the data/objective (O'neil, 2017; Mehrabi et al., 2021), affirmative action to induce longer-term changes in the macro environment (Kleinberg & Raghavan, 2018; Celis et al., 2020; 2021; Hu & Chen, 2018), or diversification as a means to safeguard against uncertainty (e.g. in a financial portfolio). This notion of diversity is often uncorrelated or inversely correlated with our understanding of utility u.

Fairness constraints based on group memberships (Binns, 2020; Dwork et al., 2012) typically enforce group-wise cardinality constraints on the chosen subset (e.g. require S to be an independent set of a partition matroid). In practice, however, when we require representation across a large number of disjoint groups or any number of intersectional underrepresented groups, a natural constraint is to enforce a minimum number (which shall be termed as **quotas**) of items chosen from each group. A common example of this is affirmative action policies such as Baswana et al. (2019) or variations of the Rooney Rule.

The contribution of this paper is an analysis of quota constraints for two popular and illustrative circumstances: maximizing a monotone submodular function and sampling from a determinantal point process. We provide algorithms with guarantees for both problems and conclude with an empirical discussion of the utilitarian price of such constraints.

2 SUBMODULAR MAXIMIZATION SUBJECT TO QUOTAS

Formally stated, let the elements of a ground set V belong to p protected groups $\{V_j \subseteq V\}_{j=1}^p$ *that may be disjoint or intersecting* with quotas $\{k_j \in \mathbb{R}\}_{j=1}^p$. We wish to choose

$$S = \arg\max_{S \in \mathcal{Q}} u(S) \quad \text{where } \mathcal{Q} = \{A : |A \cap V_j| \geq k_j \ \forall j \in 1, ..., p \text{ and } |A| \leq k\} \tag{1}$$

where u is a normalized (i.e. $u(\emptyset) = 0$), non-decreasing submodular function ($u(a|A) \leq u(a|B)\ \forall a \in V, B \subseteq A \subseteq V$ where $u(a|A)$ stands for $u(A \cup a) - u(A)$).

Note that the set $\mathcal{Q}$ of allowable subsets of V is *not* a matroid. A matroid is characterized as a set of *independent* sets $\mathcal{I} \subseteq 2^U$ that is downward closed, i.e., $A \in \mathcal{I}$ and $B \subseteq A \Rightarrow B \in \mathcal{I}$, and such that all maximal elements of $\mathcal{I}$ have the same cardinality. If we reframe the quota constraints as

$$S = \arg\max_{S \in \mathcal{Q}} u(S) \quad \text{where } \mathcal{Q} = \{A : |A \cap (V \setminus V_j)| \leq (k - k_j)\ \forall j \in 1, ..., p \text{ and } |A| \leq k\} \quad (2)$$

we have a problem equivalent to 1, which indicates that $\mathcal{Q}$ is an intersection of matroids. A caveat here is that the existence of a feasible solution for 1 (i.e. $\sum j = 1^p k_j < k$) does not assure a feasible solution for 1. The theory and applications of maximizing set functions (especially submodular or submodular-supermodular) subject to (one or many) matroidal constraints are extremely well-studied under various circumstances (Nemhauser et al., 1978; Edmonds, 1968; Calinescu et al., 2011; Do & Neumann, 2020; Friedrich et al., 2019; Buchbinder et al., 2019). A regular greedy heuristic (Algorithm 2) for this is presented in Appendix A. However, phrasing the problem as 1 allows us an alternative approach – Algorithm 1. In the QUOTAGREEDY heuristic, we iteratively build S in two stages – first focusing on meeting the quota constraint, followed by regular greedy addition of elements.

In Appendix B, we show that **Algorithm 1's approximation guarantee under a quota constraint can be just as good as that of the regular greedy heuristic under a partition matroid constraint**, for the case where the groups $\{V_j\}_{j=1}^p$ are a partition of V (i.e. are disjoint and exhaustive).

3 SAMPLING FROM A QUOTA CONSTRAINED DPP

Let us now extend the concept from Celis et al. (2018) to the quota-constrained setting, in order to understand the probability space of subsets satisfying geometric diversity (with more diverse sets being more likely to be sampled) as well as group-wise diversity (satisfied for every sample). Instead of restricting group memberships to follow exact capacity constraints, we define the problem as: sample S according to the distribution

$$P(S) \propto \begin{cases} \det(X_S X_S^\top) & \text{if } S \in \mathcal{Q} = \{A : |A \cap V_j| \geq k_j\ \forall j \in 1, ..., p \text{ and } |A| \leq k\} \\ 0 & \text{otherwise} \end{cases} \quad (3)$$

where $X \in \mathbb{R}^{|V| \times m}$ is a feature matrix characterizing the ground set. We can slightly modify their linear-time Sample-And-Project algorithm to extend to the Quota-constrained case (see Appendix C) and still expect a very analogous performance guarantee to hold under the conditions specified in the original paper. Moreover, since a quota constraint is less rigid than an exact capacity constraint, we also expect the *price of fairness* as defined in Celis et al. (2018) – the KL-divergence $D_{\mathrm{KL}}(q \| r)$ (where q is the distribution defined by 3 and r is the unconstrained distribution over $\{A : |A| \leq k\}$) to be upper bounded by the exact capacity counterpart.

4 DISCUSSION

The motivation to analyze quota constraints arose from the observation that often, members belonging to underrepresented groups often contribute disproportionately more toward the utility of the cohort as a whole when chosen alongside more members of the same group, as opposed to chosen in isolation (e.g. "token diversity hires" or isolated samples in data summaries which lack context). While a formal analysis along the lines of modeling a "true" utility that takes in-group peer effects into account is relegated for future work, an attempt is made to demonstrate that quota-like constraints can be employed alongside tweaks to existing algorithms without incurring an extra *price of fairness*, i.e. without compromising on the attained utility of the subset as a whole, as compared to group capacity constraints. Another advantage of QUOTAGREEDY not captured in the current formalism is that filling up underrepresented group quotas with high-leverage members first (as opposed to when the gains to u are diminished after reaching capacity on the overrepresented groups) will result in higher quality representation across said groups. In Appendix D, we try out the CAPACITYGREEDY and QUOTAGREEDY algorithms on realistically constructed synthetic data and observe that the latter results in much better group-wise diversity while incurring a very negligible hit to the utility.

URM STATEMENT

We acknowledge that the key author of this work meets the URM criteria of ICLR 2023 Tiny Papers Track. Additionally, they are not affiliated with a funded organization or team whose primary goal is research.

REFERENCES

Surender Baswana, Partha Pratim Chakrabarti, Sharat Chandran, Yashodhan Kanoria, and Utkarsh Patange. Centralized admissions for engineering colleges in india. In *Proceedings of the 2019 ACM Conference on Economics and Computation*, pp. 323–324, 2019.

Reuben Binns. On the apparent conflict between individual and group fairness. In *Proceedings of the 2020 conference on fairness, accountability, and transparency*, pp. 514–524, 2020.

Niv Buchbinder, Moran Feldman, and Mohit Garg. Deterministic $(1/2+\varepsilon)$-approximation for submodular maximization over a matroid. In *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms*, pp. 241–254. SIAM, 2019.

Gruia Calinescu, Chandra Chekuri, Martin Pal, and Jan Vondrák. Maximizing a monotone submodular function subject to a matroid constraint. *SIAM Journal on Computing*, 40(6):1740–1766, 2011.

L Elisa Celis, Amit Deshpande, Tarun Kathuria, and Nisheeth K Vishnoi. How to be fair and diverse? *arXiv preprint arXiv:1610.07183*, 2016.

L Elisa Celis, Vijay Keswani, Damian Straszak, Amit Deshpande, Tarun Kathuria, and Nisheeth K Vishnoi. Fair and diverse dpp-based data summarization. *arXiv preprint arXiv:1802.04023*, 2018.

L Elisa Celis, Anay Mehrotra, and Nisheeth K Vishnoi. Interventions for ranking in the presence of implicit bias. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, pp. 369–380, 2020.

L Elisa Celis, Chris Hays, Anay Mehrotra, and Nisheeth K Vishnoi. The effect of the rooney rule on implicit bias in the long term. In *Proceedings of the 2021 ACM conference on fairness, accountability, and transparency*, pp. 678–689, 2021.

Wei Chen, Shang-Hua Teng, and Hanrui Zhang. Capturing complementarity in set functions by going beyond submodularity/subadditivity. *arXiv preprint arXiv:1805.04436*, 2018.

Anh Viet Do and Frank Neumann. Maximizing submodular or monotone functions under partition matroid constraints by multi-objective evolutionary algorithms. In *Parallel Problem Solving from Nature–PPSN XVI: 16th International Conference, PPSN 2020, Leiden, The Netherlands, September 5-9, 2020, Proceedings, Part II 16*, pp. 588–603. Springer, 2020.

Cynthia Dwork, Moritz Hardt, Toniann Pitassi, Omer Reingold, and Richard Zemel. Fairness through awareness. In *Proceedings of the 3rd innovations in theoretical computer science conference*, pp. 214–226, 2012.

Jack Edmonds. Matroid partition. *50 Years of Integer Programming 1958–2008*, pp. 199, 1968.

Uriel Feige and Rani Izsak. Welfare maximization and the supermodular degree. In *Proceedings of the 4th conference on Innovations in Theoretical Computer Science*, pp. 247–256, 2013.

Tobias Friedrich, Andreas Göbel, Frank Neumann, Francesco Quinzan, and Ralf Rothenberger. Greedy maximization of functions with bounded curvature under partition matroid constraints. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pp. 2272–2279, 2019.

Zhinoos Razavi Hesabi, Zahir Tari, Andrzej Goscinski, Adil Fahad, Ibrahim Khalil, and Carlos Queiroz. Data summarization techniques for big data—a survey. *Handbook on Data Centers*, pp. 1109–1152, 2015.

Lily Hu and Yiling Chen. A short-term intervention for long-term fairness in the labor market. In *Proceedings of the 2018 World Wide Web Conference*, pp. 1389–1398, 2018.

Jon Kleinberg and Manish Raghavan. Selection problems in the presence of implicit bias. *arXiv preprint arXiv:1801.03533*, 2018.

Andreas Krause and Daniel Golovin. Submodular function maximization. *Tractability*, 3:71–104, 2014.

Hui Lin and Jeff Bilmes. A class of submodular functions for document summarization. In *Proceedings of the 49th annual meeting of the association for computational linguistics: human language technologies*, pp. 510–520, 2011.

Ninareh Mehrabi, Fred Morstatter, Nripsuta Saxena, Kristina Lerman, and Aram Galstyan. A survey on bias and fairness in machine learning. *ACM Computing Surveys (CSUR)*, 54(6):1–35, 2021.

George L Nemhauser, Laurence A Wolsey, and Marshall L Fisher. An analysis of approximations for maximizing submodular set functions—i. *Mathematical programming*, 14(1):265–294, 1978.

Cathy O'neil. *Weapons of math destruction: How big data increases inequality and threatens democracy*. Crown, 2017.

A APPENDIX: THE QUOTAGREEDY ALGORITHM

Algorithm 1: QUOTAGREEDY

Input: Universal set V,
membership matrix $\mu \in \{0,1\}^{|V|\times p}$ over p groups,
group quotas $\{k_j\}_{j=1}^{p}$, total set capacity k,
oracle access to the marginal utility gains $u(a|A)\ \forall a \in V, A \subseteq V$

Initialize $S \leftarrow \emptyset$, $u(S) \leftarrow u(\emptyset)$;
// Quota-filling stage
$\mathcal{R} \leftarrow \{j : \sum_{v\in S} \mu_{vj} < k_j\}$
while $|\mathcal{R}| > 0$ **do**
  // Search over all inadequately-represented groups $\mathcal{R}$
  $S \leftarrow S \cup \underset{\substack{v\in V\setminus S \\ \sum_{j\in\mathcal{R}} \mu_{vj}>0}}{\arg\max}\ u(v|S)$
  $\mathcal{R} \leftarrow \{j : \sum_{v\in S} \mu_{vj} < k_j\}$
end
// Regular greedy stage
for $t = |S|+1, ..., k$, **do**
  // Add the best elements
  $S \leftarrow S \cup \underset{v\in V\setminus S}{\arg\max}\ u(v|S)$
end

Algorithm 2: CAPACITYGREEDY

Input: Universal set V,
membership matrix $\mu \in \{0,1\}^{|V|\times p}$ over p groups,
group quotas $\{k_j\}_{j=1}^{p}$, total set capacity k,
oracle access to the marginal utility gains $u(a|A)\ \forall a \in V, A \subseteq V$

Initialize $S \leftarrow \emptyset$, $u(S) \leftarrow u(\emptyset)$;
// Having a minimum quota for a certain group is equivalent to
// having a maximum capacity for elements not belonging to that group
Set $h_j = k - k_j\ \forall j = 1, ..., p$;
$\mathcal{R} \leftarrow \{j : \sum_{v\in S}(1-\mu_{vj}) == h_j\}$
for $t = 1, ..., k$, **do**
  // Search over all elements which are
  // NOT members of groups $\mathcal{R}$ that have reached capacity
  $S \leftarrow S \cup \underset{\substack{v\in V\setminus S \\ \sum_{j\in\mathcal{R}}(1-\mu_{vj})==0}}{\arg\max}\ u(v|S)$
  $\mathcal{R} \leftarrow \{j : \sum_{v\in S}(1-\mu_{vj}) == h_j\}$
end

B APPENDIX: A GUARANTEE FOR QUOTAGREEDY FOR SUBMODULAR FUNCTIONS

To aid us in our endeavor, we will use as a starting point, any appropriate result for the partition matroid version of our problem statement. (Note: this problem is different from 1!) Here, $\{V_j\}_{j=1}^{p}$ are disjoint and each of them is to contain exactly k_j elements.

$$S = \underset{S\in\mathcal{Q}}{\arg\max}\ u(S) \quad \text{where } \mathcal{Q} = \{A : |A \cap V_j| = k_j\ \forall j \in 1, ..., p\} \tag{4}$$

Lemma B.1 (**Approximation Guarantee for Greedy with a Partition Matroid Constraint**). *If $S_m^\dagger$ is the optimal value of 4 and $\widetilde{S_m}$ is the solution of greedy on 4, then $u(\widetilde{S_m}) \geq \alpha\ u(S_m^\dagger)$.*

Nemhauser et al. (1978) have shown α to be $1/2$ and Calinescu et al. (2011) $\left(1-\frac{1}{e}\right)$ with a randomized method, but we shall continue to denote the approximation guarantee as a variable α so as to allow substituting extensions to the various special cases (e.g. curvature bounds, different oracle models, etc.).

Note that when $\sum_{j=1}^{p} k_j < k$, i.e. when Problem 1 has a feasible solution, and when u is normalized and non-decreasing, any feasible solution of 4 is also a feasible solution of 1.

Theorem B.2 (**Approximation Guarantee for Algorithm 1**). *If $S^\dagger$ is the true optimum of Problem 1 for a submodular, non-negative, normalized utility u and for disjoint subsets V_j, and $\widetilde{S}$ is the output of Algorithm 1, then $u(\widetilde{S}) \geq \alpha\, u(S^\dagger)$, where α is the value from Lemma B.1.*

Proof.

1. We can always arbitrarily divide $S^\dagger$ into two sets $S_Q^\dagger$ and $S_R^\dagger$ such that $S_Q^\dagger$ satisfies all constraints in 1 (and hence also all constraints in 4) with equality.
$$S^\dagger = S_Q^\dagger \sqcup S_R^\dagger \tag{5}$$
2. Observe that when the groups V_j are disjoint and exhaustive, the "quota-filling stage" in Algorithm 1 is equivalent to regular greedy under a partition matroid constraint. Thus,
$$\frac{1}{\alpha} u(\widetilde{S_{\text{m}}}) \geq u(S_{\text{m}}^\dagger) \tag{6}$$
$$\geq u(S_Q^\dagger) \;\; \text{(from the optimality of } u(S_{\text{m}}^\dagger)) \tag{7}$$
3. Submodularity and non-negativity imply subadditivity.
$$u(S^\dagger) \leq u(S_Q^\dagger) + u(S_R^\dagger) \tag{8}$$
$$\leq \frac{1}{\alpha} u(\widetilde{S_{\text{m}}}) + u(S_R^\dagger) \tag{9}$$
$$\leq \frac{1}{\alpha} u(\widetilde{S}) + u(\widetilde{S}_m \cup S_R^\dagger) \;\; (u \text{ is monotone}) \tag{10}$$
4. Regular greedy for the remaining capacity implies
$$[u(\widetilde{S}) - u(\widetilde{S}_{\text{m}})] = u(\widetilde{S}|\widetilde{S}_{\text{m}}) \tag{11}$$
$$\geq \left(1-\frac{1}{e}\right) \max_{S:|S|\leq k} u(S|\widetilde{S}_{\text{m}}) \tag{12}$$
$$\geq \left(1-\frac{1}{e}\right) [u(\widetilde{S}_{\text{m}} \cup S_R^\dagger) - u(\widetilde{S}_{\text{m}})] \tag{13}$$

   Thus, 10 and 13 are combined to give
$$u(S^\dagger) \leq \frac{1}{\alpha} u(\widetilde{S}) + \left(\frac{e}{e-1}\right) [u(\widetilde{S}) - u(\widetilde{S}_{\text{m}})] + u(\widetilde{S}_{\text{m}}) \tag{14}$$
$$= \left(\frac{1}{\alpha}\right) u(\widetilde{S}) + \left(\frac{e}{e-1}\right) u(\widetilde{S}) + \left(\frac{-1}{e-1}\right) u(\widetilde{S}_{\text{m}}) \tag{15}$$
$$\leq \left(\frac{1}{\alpha}\right) u(\widetilde{S}) + \left(\frac{e}{e-1}\right) u(\widetilde{S}) + \left(\frac{-1}{e-1}\right) u(\widetilde{S}) \tag{16}$$
$$= \left(\frac{1}{\alpha}\right) u(\widetilde{S}) \tag{17}$$

□

We have hereby shown that QUOTAGREEDY has at least as good of a guarantee as regular greedy under a partition matroid constraint. Here, we make an important note that CAPACITYGREEDY does not always output a feasible solution to Problem 1 (see Appendix D for an example).

C APPENDIX: QUOTA-AWARE-SAMPLE-AND-PROJECT

Algorithm 3: QUOTASAMPLEANDPROJECT

Input: Universal set V,
membership matrix $\mu \in \{0,1\}^{|V|\times p}$ over p groups,
group quotas $\{k_j\}_{j=1}^{p}$, total set capacity k,
feature matrix $X \in \mathbb{R}^{|V|\times m}$

Initialize $w_v = X_v \in \mathbb{R}^m \; \forall v \in V$

// Quota-filling stage

$\mathcal{R} \leftarrow \left\{ j : \sum_{v\in S} \mu_{vj} < k_j \right\}$

while $|\mathcal{R}| > 0$ **do**

  // Sample from members belonging to
  // inadequately-represented partitions

  Sample $\widetilde{v}$ from distribution $\left\{ \frac{||w_{\widetilde{v}}||^2}{\sum_{v\in R} ||w_v||^2} \right\}_{\substack{\widetilde{v}\in V\setminus S \\ \sum_{j\in\mathcal{R}} \mu_{\widetilde{v}j}>0}}$

  $S \leftarrow S \cup \widetilde{v}$

  // Project all feature vectors onto
  // the subspace orthogonal to $\widetilde{v}$

  Set $w_v \leftarrow \pi_{\widetilde{v}}(w_v) \;\; \forall v \in V \setminus S$

  // Recompute inadequately-represented partitions

  $\mathcal{R} \leftarrow \left\{ j : \sum_{v\in S} \mu_{vj} < k_j \right\}$

end

// Regular sample-and-project stage

for $t = |S|+1, ..., k$, **do**

  // Sample from all remaining elements

  Sample $\widetilde{v}$ from distribution $\left\{ \frac{||w_{\widetilde{v}}||^2}{\sum_{v\in R} ||w_v||^2} \right\}_{\widetilde{v}\in V\setminus S}$

  $S \leftarrow S \cup \widetilde{v}$

  Set $w_v \leftarrow \pi_{\widetilde{v}}(w_v) \;\; \forall v \in V \setminus S$

end

D APPENDIX: EXPERIMENTS ON SYNTHETIC DATA

In order to get a feel for how QUOTAGREEDY would play out in practice, we simulate it on some synthetic data. We used a ground set of size $|V| = 100$, selection budget of $k = 20$, and a simple mixture of concave-over-modular utilities $u(A) = \sum_{i=1}^{m} w_i \sqrt{\sum_{a\in A} X_{ai}}$ where $X \in \mathbb{R}_+^{|V|\times m}$ is a randomly drawn feature matrix, w_i are randomly instantiated weights, and $m = 80$.

Observe how $u(A)$ increases monotonically with each feature X_{ai} from the vector X_a if $a \in A$. Hence, to simulate $p = 3$ groups that are underprivileged w.r.t. their value assessed by u, we randomly assigned members to groups with some probabilities, and shrunk some of their features by a factor $0 < \beta_p \leq 1$ by setting $X_{ai} \leftarrow \beta_p \, X_{ai}$ for all $a \in V_p$ and $i \in$ a randomly chosen subset of features.

In the simulation results to follow, we take representative examples of a few configurations of $\{V_j\}_{i=j}^{p}$ (group membership distributions in light brown), and select subsets using the following 4 methods:
(final membership distributions plotted as histograms; objective values $u(S)$ in plot subtitles)

1. **[red]** A random subset selection
2. **[teal]** Greedy maximization without any group-wise constraints (i.e. over $\{A : |A| \leq k\}$)
3. **[yellow]** CAPACITYGREEDY 2 maximization with quota constraints expressed as capacity constraints (i.e. over $\{A : |A \cap (V \setminus V_j)| \leq (k - k_j) \; \forall j \in 1, ..., p$ and $|A| \leq k\}$)
4. **[orange]** QUOTAGREEDY 1 with group-wise quota constraints (over $\{A : |A \cap V_j| \geq k_j \; \forall j \in 1, ..., p$ and $|A| \leq k\}$)

D.1 DISJOINT GROUPS: 1

In this instance, we chose population distributions over 3 **disjoint** groups $\sim \{0.4, 0.2, 0.4\}$ with feature disadvantage vectors $||\beta_2|| > ||\beta_1|| > ||\beta_0|| = 0 \in (0,1]^m$.
Quota constraints were set as $k_1 = k_2 = k_3 = 4$.

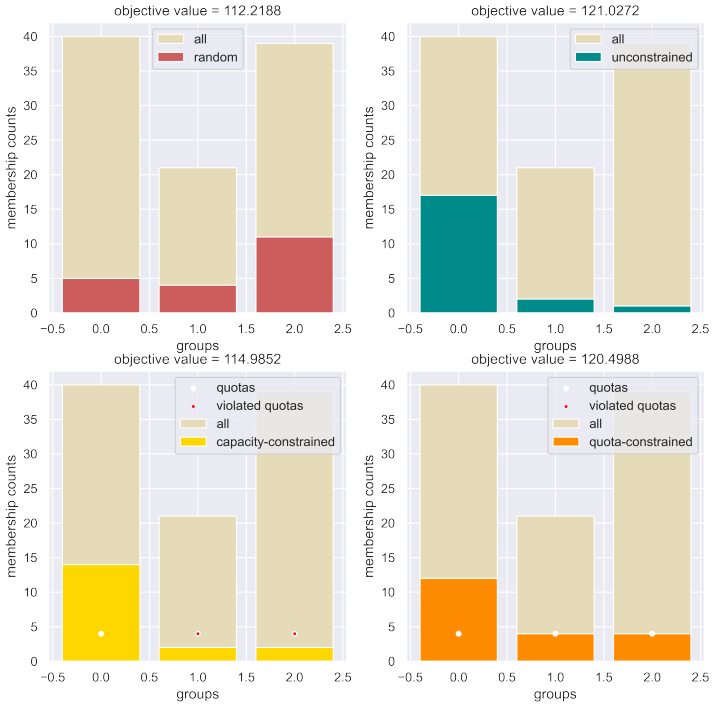

We see an example of the fact that phrasing the problem as in Algorithm 2 and prioritizing the filling of high-gain elements without hitting the quota-sensitive groups may land us in a situation where the set capacity is maxed out (it is not possible to further add elements), we are still satisfying $|A \cap (V \setminus V_j)| \leq (k - k_j)\ \forall j$, but not satisfying $|A \cap V_j| \geq k_j\ \forall i$. The objective attained by QUOTAGREEDY is almost as good as that attained without quota constraints, and much better than a random selection, pointing to a low price of fairness.

D.2 DISJOINT GROUPS: 2

This instance is identical to the one above, with a reduction in the magnitudes of feature disadvantage vectors $||\beta_p||$.

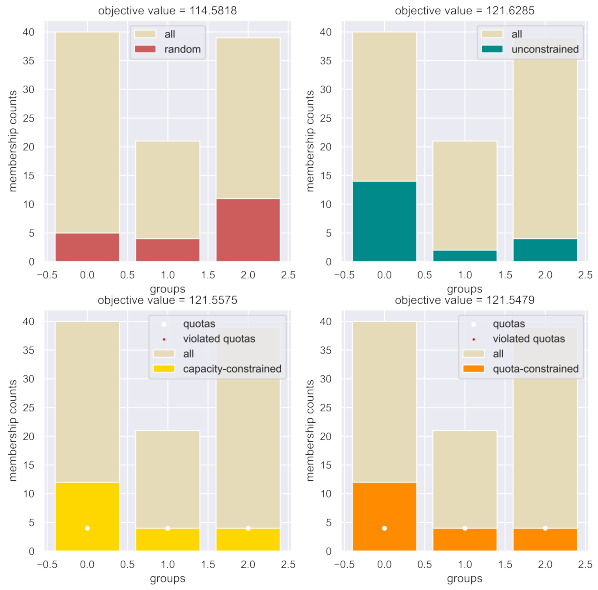

Here, CAPACITYGREEDY outperforms QUOTAGREEDY by a tiny bit, but the difference is still negligible given that both are near the unconstrained max.

D.3 INTERSECTING GROUPS

In this instance, we sample group assignments over 3 **intersecting** groups with independent probabilities $\sim \{0.4, 0.3, 0.6\}$ with feature disadvantage vectors $||\beta_2|| > ||\beta_1|| > ||\beta_0|| \in (0, 1]^m$. Quota constraints were set as $k_1 = k_2 = k_3 = 4$.

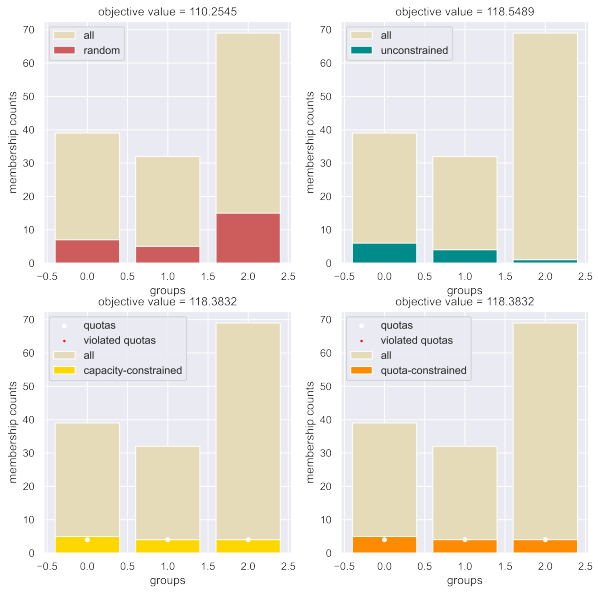

Here too, both algorithms with constraints perform much better than random, and almost as well as group-agnostic greedy.

In the intersecting group paradigm, it is also not rare to see CAPACITYGREEDY and QUOTAGREEDY render different distributions over groups – where the choice of more intersectional members in QUOTAGREEDY result in a distribution with more underrepresented members overall.

[For the sake of reproducibility, a link to this code will be provided upon request and included in the final version of this paper.]