

Variational Fair Clustering

Imtiaz Masud Ziko¹, Jing Yuan², Eric Granger¹ and Ismail Ben Ayed¹

¹ ÉTS Montreal, Canada

² Xidian University, China

imtiaz-masud.ziko.1@etsmtl.ca

Abstract

We propose a general variational framework of fair clustering, which integrates an original Kullback-Leibler (KL) fairness term with a large class of clustering objectives, including prototype or graph based. Fundamentally different from the existing combinatorial and spectral solutions, our variational multi-term approach enables to control the trade-off levels between the fairness and clustering objectives. We derive a general tight upper bound based on a concave-convex decomposition of our fairness term, its Lipschitz-gradient property and the Pinsker’s inequality. Our tight upper bound can be jointly optimized with various clustering objectives, while yielding a scalable solution, with convergence guarantee. Interestingly, at each iteration, it performs an independent update for each assignment variable. Therefore, it can be easily distributed for large-scale datasets. This scalability is important as it enables to explore different trade-off levels between the fairness and clustering objectives. Unlike spectral relaxation, our formulation does not require computing its eigenvalue decomposition. We report comprehensive evaluations and comparisons with state-of-the-art methods over various fair clustering benchmarks, which show that our variational formulation can yield highly competitive solutions in terms of fairness and clustering objectives.

Introduction

Machine learning models are impacting our daily life, for instance, in marketing, finance, education, and even in sentencing recommendations (Kleinberg et al. 2017). However, these models may exhibit biases towards specific demographic groups due to, for instance, the biases that exist within the data. For example, a higher level of face recognition accuracy may be found with white males (Buolamwini and Gebru 2018). These biases have recently triggered substantial interest in designing fair algorithms for the supervised learning setting (Hardt, Price, and Srebro 2016; Zafar et al. 2017; Donini et al. 2018). Also, very recently, the community has started to investigate fairness constraints in unsupervised learning (Chierichetti et al. 2017; Kleindessner et al. 2019; Backurs et al. 2019; Samadi et al. 2018; Celis et al. 2018). Specifically, Chierichetti et al. (Chierichetti et al. 2017) pioneered the concept of *fair clustering*. The problem consists

of embedding fairness constraints that encourage clusters to have balanced demographic groups pertaining to some sensitive attributes (e.g., sex, gender, race, etc.), so as to counteract any form of data-inherent bias.

Assume that we are given N data points to be assigned to a set of K clusters, and let $S_k \in \{0, 1\}^N$ denotes a binary indicator vector whose components take value 1 when the point is within cluster k , and 0 otherwise. Also suppose that the data contains J different demographic groups, with $V_j \in \{0, 1\}^N$ denoting a binary indicator vector of demographic group j . The authors of (Chierichetti et al. 2017; Kleindessner et al. 2019) suggested to evaluate fairness in terms of cluster-balance measures, which take the following form:

$$\text{balance}(S_k) = \min_{j \neq j'} \frac{V_j^t S_k}{V_{j'}^t S_k} \in [0, 1] \quad (1)$$

The higher this measure, the fairer the cluster. The overall clustering balance is defined by the minimum of Eq. (1) over k . This notion of fairness in clusters has recently given rise to a new line of research that was introduced, mostly, for prototype-based clustering, e.g., K-center, K-medians and K-means (Chierichetti et al. 2017; Backurs et al. 2019; Schmidt, Schwiegelshohn, and Sohler 2018; Bera et al. 2019). Also, very recently, fairness has been investigated in the context of spectral graph clustering (Kleindessner et al. 2019). The general problem raises several interesting questions. How to embed fairness in popular clustering objectives? Can we control the trade-off between some “acceptable” fairness level (or tolerance) and the quality of the clustering objective? What is the cost of fairness with respect to the clustering objective and computational complexity?

Chierichetti et al. (Chierichetti et al. 2017) investigated combinatorial approximation algorithms for maximizing the the fairness measures in Eq. (1), for K-center and K-medians clustering, and for binary demographic groups ($J = 2$). They compute *fairlets*, which are groups of points that are fair, and can not be split further into more subsets that are also fair. Then, they consider each fairlet as a data point, and cluster them with approximate K-center or K-medians algorithms. Unfortunately, as reported in the experiments in (Chierichetti et al. 2017), obtaining fair solutions with these fairlets-based algorithms comes at the price of a substantial increase in the clustering objectives. Also, the cost for finding fairlets with perfect matching is quadratic w.r.t the number of data points,

a complexity that increases for more than two demographic groups. Several combinatorial solutions followed-up on the work in (Chierichetti et al. 2017) to reduce this complexity. For instance, Backurs et al. (Backurs et al. 2019) proposed a solution to make the fairlet decomposition in (Chierichetti et al. 2017) scalable for $J = 2$, by embedding the input points in a tree metric. Rösner and Schmidt (Rösner and Schmidt 2018) designed a 14-approximate algorithm for fair K-center. (Schmidt, Schwiegelshohn, and Sohler 2018; Huang, Jiang, and Vishnoi 2019) proposed fair K-means/K-medians based on coresets – a reduced proxy set for the full dataset. Bera et al. (Bera et al. 2019) provided a bi-criteria approximation algorithm for fair prototype-based clustering, enabling multiple groups ($J > 2$). It is worth noting that, for large-scale data sets, (Chierichetti et al. 2017; Rösner and Schmidt 2018; Bera et al. 2019) sub-sample the inputs to mitigate the quadratic complexity w.r.t N . More importantly, the combinatorial algorithms discussed above are tailored for specific prototype-based objectives. For instance, they are not applicable to the very popular graph-clustering objectives, e.g., Ratio Cut or Normalized Cut (Von Luxburg 2007), which limits applicability in a breadth of graph problems, in which data takes the form of pairwise affinities.

Kleindessner et al. (Kleindessner et al. 2019) integrated fairness into graph-clustering objectives. They embedded linear constraints on the assignment matrix in spectral relaxation. Then, they solved a constrained trace optimization via finding the K smallest eigenvalues of some transformed Laplacian matrix. However, it is well-known that spectral relaxation has heavy time and memory loads since it requires storing an $N \times N$ affinity matrix and computing its eigenvalue decomposition – the complexity is cubic w.r.t N for a straightforward implementation, and super-quadratic for fast implementations (Tian et al. 2014). In the general context of spectral relaxation and graph partitioning, issues related to computational scalability for large-scale problems is driving an active line of recent work (Shaham et al. 2018; Ziko, Granger, and Ayed 2018; Vladymyrov and Carreira-Perpiñán 2016).

The existing fair clustering algorithms, such as the combinatorial or spectral solutions discussed above, do not have mechanisms that control the trade-off levels between the fairness and clustering objectives. Also, they are tailored either to prototype-based (Backurs et al. 2019; Bera et al. 2019; Chierichetti et al. 2017; Schmidt, Schwiegelshohn, and Sohler 2018) or graph-based objectives (Kleindessner et al. 2019). Finally, for a breadth of problems of wide interest, such as pairwise graph data, the computation and memory loads may become an issue for large-scale data sets.

Contributions: We propose a general, variational and bound-optimization framework of fair clustering, which integrates an original Kullback-Leibler (KL) fairness term with a large class of clustering objectives, including both prototype-based (e.g., K-means/K-medians) and graph-based (e.g., Normalized Cut or Ratio Cut). Fundamentally different from the existing combinatorial and spectral solutions, our variational multi-term approach enables to control the trade-off levels between the fairness and clustering objectives. We derive a general tight upper bound based on a concave-convex

decomposition of our fairness term, its Lipschitz-gradient property and the Pinsker’s inequality. Our tight upper bound can be jointly optimized with various clustering objectives, while yielding a scalable solution, with convergence guarantee. Interestingly, at each iteration, our general variational fair-clustering algorithm performs an independent update for each assignment variable. Therefore, it can easily be distributed for large-scale datasets. This scalability is important as it enables to explore different trade-off levels between fairness and the clustering objective. Unlike the constrained spectral relaxation in (Kleindessner et al. 2019), our formulation does not require computing its eigenvalue decomposition. We report comprehensive evaluations and comparisons with state-of-the-art methods over various fair-clustering benchmarks, which show that our variational method can yield highly competitive solutions in terms of fairness and clustering objectives, while being scalable and flexible.

Proposed Formulation

Let $\mathbf{X} = \{\mathbf{x}_p \in \mathbb{R}^M, p = 1, \dots, N\}$ denote a set of N data points to be assigned to K clusters, and \mathbf{S} is a soft cluster-assignment vector: $\mathbf{S} = [s_1, \dots, s_N] \in [0, 1]^{NK}$. For each point p , $s_p = [s_{p,k}] \in [0, 1]^K$ is the probability simplex vector verifying $\sum_k s_{p,k} = 1$. Suppose that the data set contains J different demographic groups, with vector $V_j = [v_{j,p}] \in \{0, 1\}^N$ indicating point assignment to group j : $v_{p,j} = 1$ if data point p is in group j and 0 otherwise. We propose the following general variational formulation for optimizing any clustering objective $\mathcal{F}(\mathbf{S})$ with a fairness penalty, while constraining each s_p within the K -dimensional probability simplex $\nabla_K = \{\mathbf{y} \in [0, 1]^K \mid \mathbf{1}^t \mathbf{y} = 1\}$:

$$\min_{\mathbf{S}} \mathcal{F}(\mathbf{S}) + \lambda \sum_k \mathcal{D}_{\text{KL}}(U || P_k) \quad \text{s.t.} \quad s_p \in \nabla_K \quad \forall p \quad (2)$$

$\mathcal{D}_{\text{KL}}(U || P_k)$ denotes the Kullback-Leibler (KL) divergence between the given (required) demographic proportions $U = [\mu_j]$ and the marginal probabilities of the demographics within cluster k :

$$P_k = [P(j|k)]; \quad P(j|k) = \frac{V_j^t S_k}{\mathbf{1}^t S_k} \quad \forall j, \quad (3)$$

where $S_k = [s_{p,k}] \in [0, 1]^N$ is the N -dimensional vector¹ containing point assignments to cluster k , and t denotes the transpose operator. Notice that, at the vertices of the simplex (i.e., for hard binary assignments), $V_j^t S_k$ counts the number of points within the intersection of demographic j and cluster k , whereas $\mathbf{1}^t S_k$ is the total number of points within cluster k .

Parameter λ controls the trade-off between the clustering objective and fairness penalty. The problem in (2) is challenging due to the ratios of summations in the fairness penalty and the simplex constraints. Expanding KL term $\mathcal{D}_{\text{KL}}(U || P_k)$ and

¹The set of N -dimensional vectors S_k and the set of simplex vectors s_p are two equivalent ways for representing assignment variables. However, we use S_k here for a clearer presentation of the problem, whereas, as will be clearer later, simplex vectors s_p will be more convenient in the subsequent optimization part.

discarding constant $\mu_j \log \mu_j$, our objective in (2) becomes equivalent to minimizing the following functional with respect to the relaxed assignment variables, and subject to the simplex constraints:

$$\mathcal{E}(\mathbf{S}) = \underbrace{\mathcal{F}(\mathbf{S})}_{\text{clustering}} + \lambda \underbrace{\sum_k \sum_j -\mu_j \log P(j|k)}_{\text{fairness}} \quad (4)$$

Observe that, in Eq. (4), the fairness penalty becomes a cross-entropy between the given (target) proportion U and the marginal probabilities P_k of the demographics within cluster k . Notice that our fairness penalty decomposes into convex and concave parts:

$$-\mu_j \log P(j|k) = \underbrace{\mu_j \log \mathbf{1}^t S_k}_{\text{concave}} - \underbrace{\mu_j \log V_j^t S_k}_{\text{convex}} \quad (5)$$

This enables us to derive the following tight bounds (auxiliary functions) for minimizing our general fair-clustering model in (4) using a quadratic bound and Lipschitz-gradient property of the convex part, along with Pinsker's inequality, and a first-order bound on the concave part.

Definition 1 $\mathcal{A}_i(\mathbf{S})$ is an auxiliary function of objective $\mathcal{E}(\mathbf{S})$ if it is a tight upper bound at current solution \mathbf{S}^i , i.e., it satisfies the following conditions:

$$\mathcal{E}(\mathbf{S}) \leq \mathcal{A}_i(\mathbf{S}), \forall \mathbf{S} \quad (6a)$$

$$\mathcal{E}(\mathbf{S}^i) = \mathcal{A}_i(\mathbf{S}^i) \quad (6b)$$

where i is the iteration index.

Bound optimizers, also commonly referred to as Majorize-Minimize (MM) algorithms (Zhang, Kwok, and Yeung 2007), update the current solution \mathbf{S}^i to the next by optimizing the auxiliary function:

$$\mathbf{S}^{i+1} = \arg \min_{\mathbf{S}} \mathcal{A}_i(\mathbf{S}) \quad (7)$$

When these updates correspond to the global optima of the auxiliary functions, MM procedures enjoy a strong guarantee: The original objective function $\mathcal{E}(\mathbf{S})$ does not increase at each iteration:

$$\mathcal{E}(\mathbf{S}^{i+1}) \leq \mathcal{A}_i(\mathbf{S}^{i+1}) \leq \mathcal{A}_i(\mathbf{S}^i) = \mathcal{E}(\mathbf{S}^i) \quad (8)$$

This general principle is widely used in machine learning as it transforms a difficult problem into a sequence of easier sub-problems (Zhang, Kwok, and Yeung 2007). Examples of well-known bound optimizers include concave-convex procedures (CCCP) (Yuille and Rangarajan 2001), expectation maximization (EM) algorithms and submodular-supermodular procedures (SSP) (Narasimhan and Bilmes 2005), among others. The main technical difficulty in bound optimization is how to derive an auxiliary function. In the following, we derive auxiliary functions for our general fair-clustering objectives in (4).

Proposition 1 (Bound on the fairness penalty) *Given current clustering solution \mathbf{S}^i at iteration i , we have the following auxiliary function on the fairness term in (4), up to*

additive and multiplicative constants, and for current solutions in which each demographic is represented by at least one point in each cluster (i.e., $V_j^t S_k^i \geq 1 \forall j, k$):

$$\begin{aligned} \mathcal{G}_i(\mathbf{S}) &\propto \sum_{p=1}^N \mathbf{s}_p^t (\mathbf{b}_p^i + \log \mathbf{s}_p - \log \mathbf{s}_p^i) \\ \text{with } \mathbf{b}_p^i &= [b_{p,1}^i, \dots, b_{p,K}^i] \\ b_{p,k}^i &= \frac{1}{L} \sum_j \left(\frac{\mu_j}{\mathbf{1}^t S_k^i} - \frac{\mu_j v_{j,p}}{V_j^t S_k^i} \right) \end{aligned} \quad (9)$$

where L is some positive Lipschitz-gradient constant verifying $L \leq N$.

Proof: We provide a detailed proof in the supplemental material. Here, we give the main technical ingredients for obtaining our bound. We use a quadratic bound and a Lipschitz-gradient property for the convex part, and a first-order bound on the concave part. We further bound the quadratic distances between simplex variables with the Pinsker's inequality (Csiszar and Körner 2011). This step avoids completely point-wise Lagrangian-dual projections and inner iterations for handling the simplex constraints, yielding scalable (parallel) updates, with convergence guarantee.

Proposition 2 (Bound on the clustering objective) *Given current clustering solution \mathbf{S}^i at iteration i , the auxiliary functions for several popular clustering objectives $\mathcal{F}(\mathbf{S})$ take the following general form:*

$$\mathcal{H}_i(\mathbf{S}) = \sum_{p=1}^N \mathbf{s}_p^t \mathbf{a}_p^i \quad (10)$$

where point-wise (unary) potentials \mathbf{a}_p^i are given in Table 1.

Proofs: We give detailed proofs in the supplemental material. Here, we provide the main technical aspects: For the Ncut objective, the derivation of the auxiliary function is based on the fact that, for positive semi-definite affinity matrix \mathbf{W} , the Ncut objective is concave (Tang et al. 2019). Therefore, the first-order approximation at the current solution is an auxiliary function. For the prototype-based objectives, deriving an auxiliary function follows from the observation that the optimal parameters \mathbf{c}_k , i.e., those that minimize the objective in closed-form, correspond to the sample means/medians within the clusters. These auxiliary functions correspond to the standard K-means and K-medians procedures, which can be viewed as bound optimizers (Tang et al. 2019).

Proposition 3 (Bound on the fair-clustering functional) *Given current clustering solution \mathbf{S}^i , at iteration i , and bringing back the trade-off parameter λ , we have the following auxiliary function for the general fair-clustering objective $\mathcal{E}(\mathbf{S})$ in Eq. (4):*

$$\mathcal{A}_i(\mathbf{S}) = \sum_{p=1}^N \mathbf{s}_p^t (\mathbf{a}_p^i + \lambda \mathbf{b}_p^i + \log \mathbf{s}_p - \log \mathbf{s}_p^i) \quad (11)$$

Proof: It is straightforward to check that sum of auxiliary functions, each corresponding to a term in the objective, is also an auxiliary function of the overall objective.

Notice that, at each iteration, our auxiliary function in (11) is the sum of *independent* functions, each corresponding to a single data point p . Therefore, our minimization problem in (4) can be tackled by optimizing each term over \mathbf{s}_p , subject to

Clustering	$\mathcal{F}(\mathbf{S})$	$\mathbf{a}_p^i = [a_{p,k}^i], \forall k$	Where
K-means	$\sum_N \sum_k s_{p,k} (\mathbf{x}_p - \mathbf{c}_k)^2$	$a_{p,k}^i = (\mathbf{x}_p - \mathbf{c}_k^i)^2$	$\mathbf{c}_k^i = \frac{\mathbf{X}^t S_k^i}{\mathbf{1}^t S_k^i}$
K-medians	$\sum_N \sum_k s_{p,k} d(\mathbf{x}_p, \mathbf{c}_k)$	$a_{p,k}^i = d(\mathbf{x}_p, \mathbf{c}_k^i)$	$\mathbf{c}_k^i = \arg \min_{p \neq q} d(\mathbf{x}_p, \mathbf{x}_q)$, d is a distance metric
Ncut	$K - \sum_k \frac{S_k^t \mathbf{W} S_k}{\mathbf{d}^t S_k}$	$a_{p,k}^i = d_p z_k^i - \frac{2 \sum_q w(\mathbf{x}_p, \mathbf{x}_q) s_{p,k}^i}{\mathbf{d}^t S_k^i}$	$z_k^i = \frac{(S_k^i)^t \mathbf{W} S_k^i}{(\mathbf{d}^t S_k^i)^2}$ $\mathbf{d} = [d_p]$, with $d_p = \sum_q w(\mathbf{x}_p, \mathbf{x}_q)$; $\forall p$ $\mathbf{W} = [w(\mathbf{x}_p, \mathbf{x}_q)]$ is an affinity matrix

Table 1: Auxiliary functions of several well-known clustering objectives.

the simplex constraint, and independently of the other terms, while guaranteeing convergence:

$$\min_{\mathbf{s}_p \in \nabla_K} \mathbf{s}_p^t (\mathbf{a}_p^i + \lambda \mathbf{b}_p^i + \log \mathbf{s}_p - \log \mathbf{s}_p^i), \forall p \quad (12)$$

Also, notice that, in our derived auxiliary function, we obtained a convex negative entropy barrier function $s_p \log s_p$, which comes from the convex part in our fairness penalty. This entropy term is very interesting as it avoids completely expensive projection steps and Lagrangian-dual inner iterations for the simplex constraint of each point. It yields closed-form updates for the dual variables of constraints $\mathbf{1}^t \mathbf{s}_p = 1$ and restricts the domain of each \mathbf{s}_p to non-negative values, avoiding extra dual variables for constraints $s_p \geq 0$. Interestingly, entropy-based barriers are commonly used in Bregman-proximal optimization (Yuan et al. 2017), and have well-known computational benefits when handling difficult simplex constraints (Yuan et al. 2017). However, they are not very common in the general context of clustering.

The objective in (12) is the sum of convex functions with affine simplex constraints $\mathbf{1}^t \mathbf{s}_p = 1$. As strong duality holds for the convex objective and the affine simplex constraints, the solutions of the Karush-Kuhn-Tucker (KKT) conditions minimize globally the auxiliary function. The KKT conditions yield a closed-form solution for both primal variables \mathbf{s}_p and the dual variables (Lagrange multipliers) corresponding to simplex constraints $\mathbf{1}^t \mathbf{s}_p = 1$.

$$\mathbf{s}_p^{i+1} = \frac{\mathbf{s}_p^i \exp(-(\mathbf{a}_p^i + \lambda \mathbf{b}_p^i))}{\mathbf{1}^t [\mathbf{s}_p^i \exp(-(\mathbf{a}_p^i + \lambda \mathbf{b}_p^i))]} \forall p \quad (13)$$

Notice that each closed-form update in (13) is within the simplex. We give the pseudo-code of the proposed fair-clustering in **Algorithm 1**. The algorithm can be used for any specific clustering objective, e.g., K-means or Ncut, among others, by providing the corresponding \mathbf{a}_p^i . The algorithm consists of an inner and an outer loop. The inner iterations updates \mathbf{s}_p^{i+1} using (13) until $\mathcal{A}_i(\mathbf{S})$ does not change, with the clustering term \mathbf{a}_p^i fixed from the outer loop. The outer iteration re-computes \mathbf{a}_p^i from the updated \mathbf{s}_p^{i+1} . The time complexity of each inner iteration is $\mathcal{O}(NKJ)$. Also, the updates are independent for each data p and, thus, can be efficiently computed in parallel. In the outer iteration, the time complexity of updating \mathbf{a}_p^i depends on the chosen clustering objective. For instance, for K-means, it is $\mathcal{O}(NKM)$,

and, for Ncut, it is $\mathcal{O}(N^2K)$ for full affinity matrix \mathbf{W} or much lesser for a sparse affinity matrix. Note that \mathbf{a}_p^i can be computed efficiently in parallel for all the clusters.

Convergence and monotonicity guarantees: Our variational model belongs to the family of MM procedures, whose theoretical guarantees are well-studied in the literature, e.g., (Vaida 2005). In fact, the MM principle can be viewed as a generalization of well-known expectation-maximization (EM). Therefore, in general, MM algorithms inherit the monotonicity and convergence guarantees of EM algorithms, as detailed in the theoretical discussion in (Vaida 2005). Theorem 3 in (Vaida 2005) states a condition for convergence of the general MM procedure to a local minimum: The auxiliary function has a unique global minimum, which should be obtained at each iteration when solving (7). This condition is important to guarantee, for instance, the monotonicity in (8). Our formulation satisfies this condition. In our case, the auxiliary function in (11) is strictly convex, as it is the sum of linear terms and a strictly convex term (the negative entropy), and is optimized under affine simplex constraints. Therefore, at each iteration, the closed-form solutions we obtained in (13) correspond to the unique global minimum of auxiliary function $\mathcal{A}_i(\mathbf{S})$ in (11). Our plots in Fig. 1 confirm the convergence and monotonicity of our general MM procedure for several fair-clustering objectives.

Exploring different trade-off levels via multiplier λ : Our variational multi-term formulation enables to explore several levels of trade-off between the clustering and fairness objectives via multiplier parameter λ , unlike the existing fair-clustering methods. In practice, we run in parallel our algorithm for several values of λ and choose the smallest value of λ that satisfies a pre-defined level of fairness error, i.e., $\mathcal{D}_{\text{KL}}(U||P_k) \leq \epsilon$. This is conceptually similar to standard penalty and augmented-Lagrangian approaches in constrained optimization, where the weights of the penalties² are gradually increased, until reaching a certain pre-defined precision (or duality gap) for the constraints; see Chapter 17.1 in (Nocedal and Wright 2006). The difference here is that we run independently for each λ , which can be implemented in parallel. As illustrated by the plots in Fig. 2,

²In standard constrained optimization, penalties typically take a quadratic form, unlike our method, which is based on a KL divergence penalty.

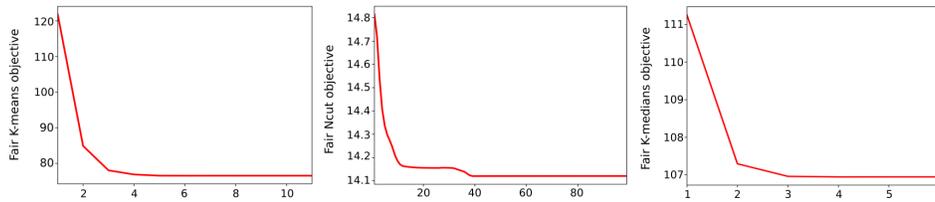


Figure 1: The convergence of the proposed bound optimizers for minimizing several fair-clustering objectives in (4): Fair K-means, Fair Ncut and Fair K- medians. The plots are based on the *Synthetic* dataset.

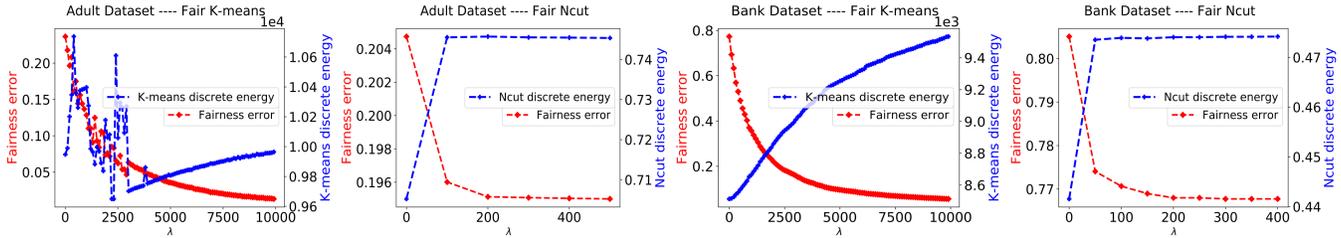


Figure 2: Clustering/Fairness objectives vs. λ .

Algorithm 1 Proposed Fair-clustering

Input: \mathbf{X} , Initial seeds, λ , U , $\{V_j\}_{j=1}^J$
Output: Clustering $labels \in \{1, \dots, K\}^N$
Initialize $labels$ from initial seeds.
Initialize \mathbf{S} from $labels$.
Initialize $i = 1$.
repeat
 Compute \mathbf{a}_p^i from \mathbf{S} (see Table 1).
 Initialize $s_p^i = \frac{\exp(-\mathbf{a}_p^i)}{\mathbf{1}^T \exp(-\mathbf{a}_p^i)}$.
 repeat
 Compute s_p^{i+1} using (13).
 $s_p^i \leftarrow s_p^{i+1}$.
 $\mathbf{S} = [s_p^i]_p$; $\forall p$.
 until $\mathcal{A}_i(\mathbf{S})$ in (11) does not change
 $i = i + 1$.
until $\mathcal{E}(\mathbf{S})$ in (4) does not change
 $l_p = \arg \max_k s_{p,k}$; $\forall p$.
 $labels = \{l_p\}_{p=1}^N$.

when λ increases, the fairness error decreases and the clustering objective increases, which is intuitive. As discussed in more details below (Tables 2, 3 and 4), our variational formulation can achieve small fairness errors (competitive with the existing state-of-the-art fair-clustering methods), but with much better clustering objectives, consistently across all the data sets.

Experiments

In this section, we present comprehensive empirical evaluations of the proposed fair-clustering algorithm, along with comparisons with state-of-the-art fair-clustering techniques.

We choose three well-known clustering objectives: K-means, K- medians and Normalized cut (Ncut), and integrate our fairness-penalty bound with the corresponding clustering bounds \mathbf{a}_p (see Table 1). We refer to our bound-optimization versions as: Fair K-means, Fair K- medians and Fair Ncut³. Note that our formulation can be used for other clustering objectives (if a bound could be derived for the objective).

We investigate the effect of fairness on the original discrete (i.e., w.r.t. binary assignment variables) clustering objectives, and compare with the existing methods. We evaluate the results in terms of the balance of each cluster S_k in (1), and define the overall **balance** of the clustering as $\text{balance} = \min_{S_k} \text{balance}(S_k)$. We further propose to evaluate the **fairness error**, which is the KL divergence $\sum_k \mathcal{D}_{\text{KL}}(U || P_k)$ in (2). This KL measure becomes equal to zero when the proportions of the demographic groups within all the output clusters match the target distribution. For Ncut, we use 20-nearest neighbor affinity matrix, \mathbf{W} : $w(\mathbf{x}_p, \mathbf{x}_q) = 1$ if data point \mathbf{x}_q is within the 20-nearest neighbors of \mathbf{x}_p , and equal to 0 otherwise. In all the experiments, we fixed $L = 2$ and found that this does not increase the objective (see the detailed explanation in the supplemental material). We standardize each dataset by making each feature attribute to have zero mean and unit variance. We then performed L2-normalization of the features, and used the standard K-means++ (Arthur and Vassilvitskii 2007) to generate initial partitions for all the models.

Datasets

Synthetic datasets: We created two types of synthetic datasets according to the proportions of the demographics, each having two clusters and a total of 400 data points in 2D features (figures in the supplemental material). The *Synthetic*

³Code is available at: <https://github.com/imtiazziko/Variational-Fair-Clustering>

Datasets	Fair K- medians			
	Objective		fairness error / balance	
	Backurs et al. 2019	Ours	Backurs et al. 2019	Ours
Synthetic ($N = 400, J = 2, \lambda = 10$)	299.859	292.4	0.00/1.00	0.00/1.00
Synthetic-unequal ($N = 400, J = 2, \lambda = 10$)	185.47	174.81	0.77/0.21	0.00/0.33
Adult ($N = 32, 561, J = 2, \lambda = 9000$)	19330.93	18467.75	0.27/0.31	0.01/0.43
Bank ($N = 41, 108, J = 3, \lambda = 9000$)	N/A	19527.08	N/A	0.02/0.18
Census II ($N = 2, 458, 285, J = 2, \lambda = 500000$)	2385997.92	1754109.46	0.41/0.38	0.02/0.78

Table 2: Comparison of the proposed Fair K- medians to (Backurs et al. 2019).

Datasets	Fair K-means			
	Objective		fairness error / balance	
	Bera et al. 2019	Ours	Bera et al. 2019	Ours
Synthetic ($N = 400, J = 2, \lambda = 10$)	758.43	207.80	0.00 / 1.00	0.00 / 1.00
Synthetic-unequal ($N = 400, J = 2, \lambda = 10$)	180.00	159.75	0.00 / 0.33	0.00 / 0.33
Adult ($N = 32, 561, J = 2, \lambda = 9000$)	10913.84	9984.01	0.018 / 0.41	0.018 / 0.41
Bank ($N = 41, 108, J = 3, \lambda = 6000$)	11331.51	9392.20	0.03 / 0.16	0.05 / 0.17
Census II ($N = 2, 458, 285, J = 2, \lambda = 500000$)	1355457.02	1018996.53	0.07/0.77	0.02/0.78

Table 3: Comparison of the proposed Fair K-means to (Bera et al. 2019).

dataset contains two perfectly balanced demographic groups, each having an equal number of 200 points. For this data set, we imposed target target proportions $U = [0.5, 0.5]$. To experiment with our fairness penalty with unequal proportions, we also used *Synthetic-unequal* dataset with 300 and 100 points within each of the two demographic groups. In this case, we imposed target proportions $U = [0.75, 0.25]$.

Real datasets: We use three datasets from the UCI machine learning repository, one large-scale data set whose demographics are balanced (Census), along with two other data sets with various demographic proportions:

*Bank*⁴ dataset contains 41188 number of records of direct marketing campaigns of a Portuguese banking institution corresponding to each client contacted (Moro, Cortez, and Rita 2014). Note that, the previous fair clustering methods (Bera et al. 2019; Backurs et al. 2019) used a much smaller version of Bank dataset with only 4520 number of records with $J = 2$ and 3 attributes. Instead, we utilize the marital status as the sensitive attribute, which contains three groups ($J = 3$) – single, married and divorced – and removed the “Unknown” marital status. Thus, we have 41, 108 records in total. We chose 6 numeric attributes (age, duration, euribor of 3 month rate, no. of employees, consumer price index and number of contacts performed during the campaign) as features. We set the number of clusters $K = 10$, and impose the target proportions of three groups $U = [0.28, 0.61, 0.11]$ within each cluster.

*Adult*⁵ is a US census record data set from 1994. The dataset contains 32, 561 records. We used the gender status as the sensitive attribute, which contains 10771 females and 21790 males. We chose the 5 numeric attributes as features, set the number of clusters to $K = 10$, and impose proportions $U = [0.33, 0.67]$ within each cluster.

*Census*⁶ is a large-scale data set corresponding to a US cen-

sus record data from 1990. The dataset contains 2, 458, 285 records. We used the gender status as the sensitive attribute, which contains 1, 191, 601 females and 1, 266, 684 males. We chose the 25 numeric attributes as features, similarly to (Backurs et al. 2019). We set the number of clusters to $K = 20$, and imposed proportions $U = [0.48, 0.52]$ within each cluster.

Results

In this section, we discuss the results of the different experiments to evaluate the proposed general variational framework for Fair K-means, Fair K- medians and Fair Ncut. We further report comparisons with (Bera et al. 2019), (Backurs et al. 2019) and (Kleindessner et al. 2019) in terms of discrete fairness measures and clustering objectives.

Trade-off between clustering and fairness objectives:

We assess the effect of imposing fairness constraints on the original clustering objectives. In each plot in Fig. 2, the blue curve depicts the discrete-valued clustering objective $\mathcal{F}(\mathbf{S})$ (K-means or Ncut) obtained at convergence as a function of λ . The fairness error is depicted in red. Observe that, when multiplier λ increases (starting from a certain value), the discrete clustering objective increases while the fairness error decreases, which is intuitive. Also, the fairness error approaches 0 when $\lambda \rightarrow +\infty$, and both the clustering and fairness objectives tend to reach a plateau starting from a certain value of λ . The scalability of our model is highly relevant because it enables us to explore several solutions, each corresponding to a different value of multiplier λ , and to choose the smallest λ (i.e., the best clustering objective) that satisfies a pre-defined fairness level $\sum_k \mathcal{D}_{\text{KL}}(U || P_k) \leq \epsilon$. As detailed below, this flexibility enabled us to obtain better solutions, in terms of fairness and clustering objectives, than several recent fair-clustering methods. Low fairness errors are typically achieved with large values of λ . This is due to the fact that the scale of the fairness penalty could be much smaller than the clustering objectives. Notice that, for relatively small

⁴<https://archive.ics.uci.edu/ml/datasets/Bank+Marketing>

⁵<https://archive.is.uci.edu/ml/datasets/adult>

⁶[https://archive.ics.uci.edu/ml/datasets/US+Census+Data+\(1990\)](https://archive.ics.uci.edu/ml/datasets/US+Census+Data+(1990))

Datasets	Fair NCUT			
	Objective		fairness error / balance	
	Kleindessner et al. 2019	Ours	Kleindessner et al. 2019	Ours
Synthetic ($N = 400, J = 2, \lambda = 10$)	0.0	0.0	0.00/1.00	0.0/1.00
Synthetic-unequal ($N = 400, J = 2, \lambda = 10$)	0.03	0.06	0.00/0.33	0.00/0.33
Adult ($N = 32, 561, J = 2, \lambda = 10$)	0.47	0.74	0.06/0.32	0.08/0.30
Bank ($N = 41, 108, J = 3, \lambda = 40$)	N/A	0.58	N/A	0.39/0.14
Census II ($N = 2, 458, 285, J = 2, \lambda = 100$)	N/A	0.52	N/A	0.41/0.43

Table 4: Comparison of the proposed Fair NCut to (Kleindessner et al. 2019).

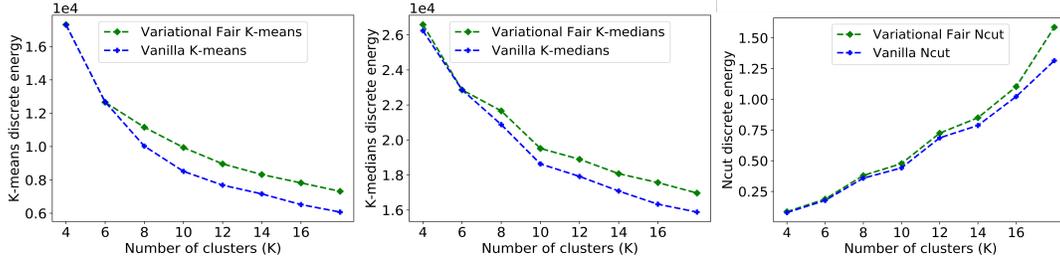


Figure 3: Effect of K on the clustering objectives for vanilla clustering and our variational fair clustering methods, including the K-means, K-medians and Ncut objectives. The results are shown for the Bank dataset. Note that, for each plot, the value of multiplier λ is fixed.

values of λ , the K-means objective (blue curve) for the Adult dataset has an oscillating behaviour. This might be due to the fact that, for small λ , the K-means objective dominates the KL fairness term. However, after a certain value of λ ($\lambda \geq 4000$), the curves become smooth, with a predictable behaviour (i.e., the fairness term decreases and the clustering term increases). When the clustering objective dominates, the oscillating behaviour might be due to the local minima of bound optimization for the K-means term. We hypothesize that, with higher values of λ , the KL fairness term “convexifies” the function, and facilitates optimization. With smaller values of λ , the K-means term dominates, with possibilities of being stuck in local minima (K-means is well-known to be sensitive to the initial conditions).

Clustering cost with respect to K : Fig. 3 depicts how the clustering objectives behave w.r.t the number of clusters K , with and without the fairness constraints. We plot the discrete clustering objective vs. K for K-means, K-medians and Ncut, using the Bank dataset, with each plot corresponding to a fixed multiplier λ . In both cases (i.e., with and without the fairness constraints), the obtained clustering objectives decrease with K , with the gap between the clustering objective obtained under fairness constraints and the vanilla clustering increasing with K . Those experimental observations are consistent with the observations in (Bera et al. 2019).

Comparisons to state-of-the-art methods: Our algorithm is flexible as it can be used in conjunction with different well-known clustering objectives. This enabled us to compare our Fair K-medians, Fair K-means and Fair Ncut versions to (Backurs et al. 2019), (Bera et al. 2019) and (Kleindessner et al. 2019), respectively. Tables 2, 3 and 4 report comparisons in terms of the original clustering objectives, achieved minimum balances and fairness errors, for Fair K-medians, Fair K-means and Fair NCut. For our model, we run the algo-

rithm for several values of λ in ascending order, and choose the smallest λ that satisfies a pre-defined level of fairness error. This flexibility and scalability enabled us to obtain significantly better clustering objectives and fairness/minimum-balance measures in comparisons to (Backurs et al. 2019); See Table 2. It is worth noting that, for the *Bank* dataset, we were unable to run (Backurs et al. 2019) as the number of demographic group is 3 (i.e. $J > 2$). In comparison to (Bera et al. 2019), our variational method achieves significantly better K-means clustering objectives, with approximately the same fairness levels. Note that, we can obtain better fairness with larger λ values. These results highlight the benefits of our proposed variational formulation, which provides control over the trade-off between the fairness level and clustering objective. In the case of fair NCut, (Kleindessner et al. 2019) achieved slightly better Ncut objectives than our model, while achieving similar fairness levels. However, we were unable to run the spectral solution of (Kleindessner et al. 2019) for large-scale *Census II* data set, and for *Bank*, due to its computational and memory load (as it requires computing the eigen values of the square affinity matrix).

Our algorithm scales up to more than two demographic groups, i.e. when $J > 2$ (e.g. *Bank*), unlike many of the existing approaches. Furthermore, for NCut graph clustering, our bound optimizer can deal with large-scale data sets, unlike (Kleindessner et al. 2019), which requires eigen decomposition for large affinity matrices. Finally, the parallel structure of our algorithm within each iteration (i.e., independent updates for each assignment variable) enables to explore different values of λ , thereby choosing the best trade-off between the clustering objective and fairness error.

Broader Impact

This paper deals with ensuring fairness criteria in clustering decisions, so as to avoid unfair treatment of minority groups pertaining to a sensitive attribute such as race, gender, etc. The paper is an endeavor to present a flexible mechanism, so as to relatively control the required fairness, while ensuring clustering quality at the same time.

References

- Arthur, D.; and Vassilvitskii, S. 2007. k-means++: The advantages of careful seeding. In *ACM-SIAM symposium on Discrete algorithms*, 1027–1035. Society for Industrial and Applied Mathematics.
- Backurs, A.; Indyk, P.; Onak, K.; Schieber, B.; Vakilian, A.; and Wagner, T. 2019. Scalable fair clustering. *International conference on machine learning (ICML)* 405–413.
- Bera, S.; Chakrabarty, D.; Flores, N.; and Negahbani, M. 2019. Fair algorithms for clustering. In *Advances in Neural Information Processing Systems*, 4955–4966.
- Buolamwini, J.; and Gebru, T. 2018. Gender shades: Intersectional accuracy disparities in commercial gender classification. In *Conference on Fairness, Accountability and Transparency*, 77–91.
- Celis, L. E.; Keswani, V.; Straszak, D.; Deshpande, A.; Kathuria, T.; and Vishnoi, N. K. 2018. Fair and Diverse DPP-Based Data Summarization. In *International Conference on Machine Learning (ICML)*, 715–724.
- Chierichetti, F.; Kumar, R.; Lattanzi, S.; and Vassilvitskii, S. 2017. Fair Clustering Through Fairlets. In *Neural Information Processing Systems (NeurIPS)*, 5036–5044.
- Csiszar, I.; and Körner, J. 2011. *Information theory: coding theorems for discrete memoryless systems*. Cambridge University Press.
- Donini, M.; Oneto, L.; Ben-David, S.; Shawe-Taylor, J.; and Pontil, M. 2018. Empirical Risk Minimization Under Fairness Constraints. In *Neural Information Processing Systems (NeurIPS)*, 2796–2806.
- Hardt, M.; Price, E.; and Srebro, N. 2016. Equality of Opportunity in Supervised Learning. In *Neural Information Processing Systems (NeurIPS)*, 3315–3323.
- Huang, L.; Jiang, S.; and Vishnoi, N. 2019. Coresets for clustering with fairness constraints. In *Advances in Neural Information Processing Systems*, 7587–7598.
- Kleinberg, J.; Lakkaraju, H.; Leskovec, J.; Ludwig, J.; and Mullainathan, S. 2017. Human decisions and machine predictions. *The quarterly journal of economics* 133(1): 237–293.
- Kleindessner, M.; Samadi, S.; Awasthi, P.; and Morgenstern, J. 2019. Guarantees for Spectral Clustering with Fairness Constraints. In *International Conference of Machine Learning (ICML)*, 3458–3467.
- Moro, S.; Cortez, P.; and Rita, P. 2014. A data-driven approach to predict the success of bank telemarketing. *Decision Support Systems* 62: 22–31.
- Narasimhan, M.; and Bilmes, J. 2005. A Submodular-supermodular Procedure with Applications to Discriminative Structure Learning. In *Conference on Uncertainty in Artificial Intelligence (UAI)*, 404–412. ISBN 0-9749039-1-4.
- Nocedal, J.; and Wright, S. 2006. *Numerical optimization*. Springer.
- Rösner, C.; and Schmidt, M. 2018. Privacy preserving clustering with constraints. In *ICALP*.
- Samadi, S.; Tantipongpipat, U. T.; Morgenstern, J. H.; Singh, M.; and Vempala, S. 2018. The Price of Fair PCA: One Extra dimension. In *Neural Information Processing Systems (NeurIPS)*, 10999–11010.
- Schmidt, M.; Schwiegelshohn, C.; and Sohler, C. 2018. Fair Coresets and Streaming Algorithms for Fair k-Means Clustering. *arXiv 1304.6478* abs/1812.10854.
- Shaham, U.; Stanton, K.; Li, H.; Basri, R.; Nadler, B.; and Kluger, Y. 2018. SpectralNet: Spectral Clustering using Deep Neural Networks. In *International Conference on Learning Representations (ICLR)*.
- Tang, M.; Marin, D.; Ayed, I. B.; and Boykov, Y. 2019. Kernel Cuts: Kernel and Spectral Clustering Meet Regularization. *International Journal of Computer Vision* 127: 477–511.
- Tian, F.; Gao, B.; Cui, Q.; Chen, E.; and Liu, T.-Y. 2014. Learning deep representations for graph clustering. In *AAAI Conference on Artificial Intelligence*, 1293–1299.
- Vaida, F. 2005. Parameter convergence for EM and MM algorithms. *Statistica Sinica* 15: 831–840.
- Vladymyrov, M.; and Carreira-Perpiñán, M. 2016. The Variational Nystrom method for large-scale spectral problems. In *International Conference on Machine Learning (ICML)*, 211–220.
- Von Luxburg, U. 2007. A tutorial on spectral clustering. *Statistics and computing* 17(4): 395–416.
- Yuan, J.; Yin, K.; Bai, Y.; Feng, X.; and Tai, X. 2017. Bregman-Proximal Augmented Lagrangian Approach to Multiphase Image Segmentation. In *Scale Space and Variational Methods in Computer Vision (SSVM)*, 524–534.
- Yuille, A. L.; and Rangarajan, A. 2001. The Concave-Convex Procedure (CCCP). In *Neural Information Processing Systems (NIPS)*, 1033–1040.
- Zafar, M. B.; Valera, I.; Gomez-Rodriguez, M.; and Gummadi, K. P. 2017. Fairness Constraints: Mechanisms for Fair Classification. In *International Conference on Artificial Intelligence and Statistics (AISTATS)*, 962–970.
- Zhang, Z.; Kwok, J. T.; and Yeung, D.-Y. 2007. Surrogate maximization/minimization algorithms and extensions. *Machine Learning* 69: 1–33.
- Ziko, I.; Granger, E.; and Ayed, I. B. 2018. Scalable Laplacian K-modes. In *Advances in Neural Information Processing Systems*, 10041–10051.