Algorithmic Greenlining: An Approach to Increase Diversity

Christian Borgs

Jennifer Chayes

Nika Haghtalab

Adam Tauman Kalai

{christian.borgs,jchayes,nika.haghtalab,adam.kalai}@microsoft.com

Microsoft Research

Cambridge, Massachusetts

Ellen Vitercik

vitercik@cs.cmu.edu

Carnegie Mellon University

Pittsburgh, Pennsylvania

ABSTRACT

In contexts such as college admissions, hiring, and image search, decision-makers often aspire to formulate selection criteria that yield both high-quality and diverse results. However, simultaneously optimizing for quality and diversity can be challenging, especially when the decision-maker does not know the true quality of any criterion and instead must rely on heuristics and intuition. We introduce an algorithmic framework that takes as input a user's selection criterion, which may yield high-quality but homogeneous results. Using an application-specific notion of substitutability, our algorithms suggest similar criteria with more diverse results, in the spirit of statistical or demographic parity. For instance, given the image search query "chairman", it suggests alternative queries which are similar but more gender-diverse, such as "chairperson". In the context of college admissions, we apply our algorithm to a dataset of students' applications and rediscover Texas's "top 10% rule": the input criterion is an ACT score cutoff, and the output is a class rank cutoff, automatically accepting the students in the top decile of their graduating class. Historically, this policy has been effective in admitting students who perform well in college and come from diverse backgrounds. We complement our empirical analysis with learning-theoretic guarantees for estimating the true diversity of any criterion based on historical data.

CCS CONCEPTS

• **Social and professional topics** → **Race and ethnicity; Gender**; • **Theory of computation** → *Design and analysis of algorithms*.

KEYWORDS

diversity, algorithms, college admissions, job applicant search, image search

ACM Reference Format:

Christian Borgs, Jennifer Chayes, Nika Haghtalab, Adam Tauman Kalai, and Ellen Vitercik. 2019. Algorithmic Greenlining: An Approach to Increase Diversity. In *AAAI/ACM Conference on AI, Ethics, and Society (AIES '19)*,

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions.acm.org.

AIES '19, January 27–28, 2019, Honolulu, HI, USA

© 2019 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-6324-2/19/01...\$15.00

<https://doi.org/10.1145/3306618.3314246>

January 27–28, 2019, Honolulu, HI, USA. ACM, New York, NY, USA, 8 pages.
<https://doi.org/10.1145/3306618.3314246>

1 INTRODUCTION

In many application domains such as college admissions, hiring, and image search, domain experts aim to develop selection criteria that yield high-quality and diverse results. However, they often operate under an unavoidable lack of information: they do not know the true quality of any given criterion, and instead must rely upon heuristics and intuition. This makes it difficult to simultaneously optimize quality and diversity. For example, consider choosing university admissions criteria, a scenario where there is extremely limited information about the true quality of any candidate. An admissions officer's intuition might suggest that admitting students with an SAT score above 1400 would provide high-quality candidates, but the criterion might admit few students from minority groups. Nonetheless, there may be similar criteria, such as "SAT score above 1200 and class rank in the top 10%" which yield high-quality and diverse student bodies. Of course, a human decision-maker cannot search over all related criteria by hand in order to discover criteria that return diverse results. In this work, we present a framework for automating this search procedure.

Our algorithms take as input a user-specified criterion t , a pre-defined similarity measure among criteria, and a set of examples. They suggest alternate criteria $t'_1, \dots, t'_k$ with greater estimated diversity on the set of examples but which are similar to t . For example, t might be the image search query "haircut" and $t'_1, \dots, t'_k$ might include related searches such as "hairstyles" and "popular haircuts". We model a criterion t as a function $t : X \rightarrow R$, where X is a set of observables, such as job applications, and R is a set of results, perhaps indicating whether an applicant is hired. Our algorithms rely on application-specific similarity functions which measure how substitutable any two criteria are. They also depend on functions that measure the diversity of any criterion's results. Our algorithms optimize the similarity function while meeting a desired diversity constraint.

We apply our framework in three key areas: college admissions, job search, and image search. In each application, we suggest

This research uses public data from the Texas Higher Education Opportunity Project (THEOP) and acknowledges the following agencies that made THEOP data available through grants and support: Ford Foundation, The Andrew W. Mellon Foundation, The William and Flora Hewlett Foundation, The Spencer Foundation, National Science Foundation (NSF Grant SES-0350990), The National Institute of Child Health & Human Development (NICHD Grant R24 H0047879) and The Office of Population Research at Princeton University.

straight-forward, natural substitutability and diversity measures, and then exemplify how our techniques can be used to find similar criteria that yield more diverse results. For example, in the case of college admissions, given two criteria t and t' and the sets E_t and $E_{t'}$ of students they admit, we measure the similarity of t and t' as a simple function of the symmetric difference between E_t and $E_{t'}$. As another example which illustrates the generality of our approach, we say that two jobs are similar if the requisite skills, education, and training are transferable, as formalized by the Department of Labor. We emphasize that in all of our applications, a domain expert could alter these two measures as they see fit. Our goal is not to nail down a universally optimal definition of substitutability and diversity, but to introduce this technique to domain experts who can determine appropriate metrics for their specific applications. We also provide provable guarantees for estimating the diversity of any criterion given historical data. Thus, our algorithms can use diversity estimates instead of true diversity scores.

We call our approach *algorithmic greenlining* as it is the antithesis of redlining, the historic and systematic denial of services to residents of specific communities, often due to demographics. Greenlining is the conscious effort to promote minority interests and representation by using unbiased criteria in day-to-day settings. In this work, we provide individuals with computational tools to help them reduce the intentional or unintentional bias encoded by their criteria.

2 RELATED WORK

Algorithmic fairness using similarity metrics. Our notion of criterion similarity is reminiscent of the similarity metric used over people for *individually fair classification* [8]. As defined by [8], under an individually fair classification system, similar individuals should be treated (i.e., classified) similarly. Thus, it relies on a similarity metric between individuals, often based on high-dimensional attribute vectors. In contrast, we rely on a notion of similarity between criteria.

Algorithmic fairness with and without quality scores. We do not assume the user knows the true quality of any given criterion or outcome. Thus, we cannot measure criterion fairness using notions such as equalized odds [11] which depend on ground truth. Rather, we can use any notion that only depends on the outcomes of the criteria (e.g., the fraction of students accepted under a given university admissions criterion from a minority group), such as statistical parity [8].

Image search. Several works analyze the presence of bias in image results from major search engines [14, 16]. Others study how to achieve diversity in image search outside the context of algorithmic fairness. Given a query, these works provide algorithms that return many different kinds of images matching that query (e.g., [15, 22]). For example, an image search for “apple” should not only return images of the fruit, but also of the computer. Some works also study how to ensure the image search results exhibit racial and gender diversity [5]. In contrast, our goal is to provide search suggestions whose results are more diverse, rather than altering the images returned by a search.

Fair ranking. In a vein similar to image search, many researchers have studied how to rank a set of items fairly [2, 6, 13, 17, 23, 26] and how to measure the fairness of an existing ranking [23, 24]. Unlike our work, these works assume the algorithm has access to ground-truth quality scores.

Query suggestions. In addition to ranking, Information Retrieval systems often provide suggestions of alternative relevant searches (e.g., [12]). The goal has generally been to help the user better meet their information seeking needs. However, to the best of our knowledge, no prior work in this domain has considered fairness or diversity.

We include additional related work in the appendix.

3 THE GENERAL MODEL AND NOTATION

In our model, there is an underlying set $X \times Z$ of examples (x, z) where $x \in X$ represents the observable attributes of an instance and $z \in Z$ represents its protected attributes. There is a set T of criteria. Each criterion $t \in T$ is a function mapping the set X of observable attributes to a set R of results. For example, in the case of college admissions, $R = \{0, 1\}$ and $t(x) = 1$ if and only if the student with observable attributes x is admitted under criterion t .

We want to measure the diversity of a criterion given an arbitrary multi-set $E \subseteq X \times Z$ of examples. In the admissions example, E could be the set of students who apply to some university on a given year. When the results are binary ($R = \{0, 1\}$) given a multi-set of examples E , we use the notation $E_t = \{(x, z) \mid (x, z) \in E, t(x) = 1\}$ to denote the subset of examples $(x, z) \in E$ such that $t(x) = 1$.

For every multi-set $E \subseteq X \times Z$ of examples and every criterion $t \in T$, there is a diversity score $\text{div}(t, E) \in [0, 1]$. For example, t might be a threshold indicating that all students with an SAT score above 600 should be admitted. Given the set E of students, $\text{div}(t, E)$ might measure the fraction of admitted students who are Hispanic.

We assume we have a substitutability function σ where $\sigma(t, t', E) \in [0, 1]$ indicates how good a substitute t' is for t given the multi-set of examples E , not accounting for fairness or diversity. This captures how willing the agent would be to replace their criterion t with the criterion t' . For instance, someone searching for “doctor” might be happy to instead search for “physician”. If each criterion is a range of admissible standardized test scores, a small change in thresholds might lead to a much more diverse set of admits. We assume that $\sigma(t, t, E) = 1$ for all $t \in T$ and all multi-sets $E \subseteq X \times Z$. We occasionally refer to the distance $\text{dist}(t, t', E)$ between two criteria (using the word “distance” loosely: we do not require that dist be symmetric or that it satisfy the triangle inequality) and we define $\text{dist}(t, t', E) = 1 - \sigma(t, t', E)$.

This definition of substitutability is extremely versatile: it can measure both intrinsic and extrinsic similarity. As an example of intrinsic similarity, we say at a high level that two jobs are similar if the requisite skills, education, and training are transferable, as formalized by the Department of Labor. In this case, $\sigma(t, t', E)$ is independent of E . Meanwhile, as an example of extrinsic similarity, we say that two university admissions criteria are similar if moving from one criterion to the other does not drastically change the set of admitted students. Thus, similarity is not intrinsic to the criteria themselves, but is a function of how they interact with the dataset at hand, i.e., $\sigma(t, t', E)$ depends on E as well as t and t' . Again, we

emphasize that the similarity metrics we employ are only examples. A domain expert could employ any definition they see fit.

Given an initial criterion t , our goal is to suggest a criterion t' that is largely substitutable ($\sigma(t, t', E)$ is close to 1) and is diverse ($\text{div}(t, E)$ is within a user-specified range).

Notation. We denote the i^{th} entry of a vector $\mathbf{v}$ as $v[i]$.

4 GENERAL MODEL INSTANTIATIONS

We now instantiate our model in several settings: university admissions, image search, and job applicant search. In the appendix, we discuss the ethics of our data usage.

4.1 University admissions

We begin with a problem motivated by university admissions. We represent a set of applicants by feature vectors in $[0, H]^d$, for some $H \in \mathbb{R}$. We use data collected from the University of Texas at Austin by the Texas Higher Education Opportunity Project (THEOP). The features we focus on are the students' composite SAT scores, composite ACT scores, and high school class rank (higher is better). We select our minority group to consist of applicants who are identified as "Black, Non-Hispanic" or "Hispanic" ($z = 1$) and our majority group to consist of "White, Non-Hispanic" applicants ($z = 0$). Therefore, $X = [0, H]^d$ and $Z = \{0, 1\}$.

In our first experiment, the set T of criteria consists of all axis-aligned hyperplanes intersecting $[0, H]^d$. Specifically, each criterion $t \in T$ is defined by a dimension $i \in [d]$ and a threshold $b \in [0, H]$. The multi-set $E_t \subseteq E$ consists of those tuples (x, z) with $x[i] \geq b$. For any criterion t such that $E_t \neq \emptyset$, we define $f_m(t, E)$ to be the fraction of individuals accepted by t who are minorities, i.e., $f_m(t, E) = \frac{1}{|E_t|} \sum_{(x, z) \in E_t} z$. We define div to measure how far this fraction is from a user-specified diversity level c , i.e., $\text{div}(t, E) = 1 - |c - f_m(t, E)|$.

We define the substitutability of two criteria t and t' based on the overlap of the students accepted under the two rules. Using $E_t \Delta E_{t'} = (E_t \cup E_{t'}) \setminus (E_t \cap E_{t'})$ to denote the symmetric difference between E_t and $E_{t'}$, we define

$$\sigma(t, t', E) = \begin{cases} 1 - \frac{|E_t \Delta E_{t'}|}{m} & \text{if } |E_t| \leq |E_{t'}| \leq 2|E_t| \\ 0 & \text{otherwise.} \end{cases}$$

By enforcing $|E_t| \leq |E_{t'}| \leq 2|E_t|$, the number of applicants admitted by the two criteria is not drastically different.

The multi-set E of feature vectors describes the first 50,000 applicants from the dataset who a) are identified as "White, Non-Hispanic," "Black, Non-Hispanic," or "Hispanic," and b) took the ACT. Of these, 24% are from the minority group. Based on this, we set the fraction c of minority students in the definition of the function div to be $c = 0.24$. Our input criterion t accepts all students whose cumulative ACT score is above 56 (the dotted line in Figure 1b), which includes 11,484 majority students and 1,298 minority students, so $\text{div}(t, E) = 0.87$.

Our goal is to find a criterion t' such that $\text{div}(t', E) \geq 0.95$ (i.e., the fraction of admitted minority students is in $[0.19, 0.29]$) and $\sigma(t, t', E)$ is as close to 1 as possible. (The threshold 0.95 could be replaced by any other bound.) This is summarized by the following

optimization problem:

$$\text{maximize } \sigma(t, t', E) \text{ subject to } \text{div}(t', E) \geq 0.95. \quad (1)$$

The new criterion's output will be different from the original criterion's output since it will be more diverse, but the change will be as small as possible. Under the optimal solution to Equation (1), students from the top 10% of their class are admitted. This criterion t' is illustrated by the solid line in Figure 1c. It admits 14,582 majority and 5,576 minority students, so $\text{div}(t', E) = 0.96$. Moreover, $\sigma(t, t', E) = 0.69$.

We thus rediscover the "top 10% rule" passed in 1997 by the Texas legislature (H.B.588), which guarantees any Texas student in the top 10% of their graduating class admission into any of Texas's public post-secondary institutions.

In the full version of the paper, we describe our experiments in a more general setting where each criterion is represented by the intersection of multiple axis-aligned half-spaces. For example, a criterion might require that an admitted applicant has an ACT score above 50 and is in the top 40% of her class. This generalization allows us to preserve a higher degree of similarity while still increasing diversity. See Figure 2 for a visualization of our experimental results.

4.2 Image search

In image search, our goal is to provide users with similar image search terms that yield more diverse results. For example, if a user begins with the Google search "haircut", the resulting images are often surprisingly male-dominated. However, we find that there are often similar queries, such as "popular haircuts", which return more images of women.

Our model applies to any search engine, but we concentrate on Google and Bing in our experiments. We define the set X of observable attributes to be the set of all images indexed by Google and $Z = \{0, 1\}$ to be the set of sensitive attributes indicating gender. In particular, given an example (x, z) , $z = 0$ if and only if x is an image of a man. We define the set of input examples E to be the entire set $X \times Z$. Given a search criterion t such as "haircut", we define $t(x)$ to be 1 if the image x is displayed in the first fifty images returned by a Google search for t and 0 otherwise, filtering Google's results so that they only return images of faces.

We adopt a simple diversity metric: given a search term t and the top fifty images returned by this search, how far from 50-50 is the gender ratio? In other words, $\text{div}(t, E) = 1 - \left| \frac{1}{2} - \frac{1}{50} \sum_{(x, z) \in E_t} z \right|$. Given an even split, $\frac{1}{50} \sum_{(x, z) \in E_t} z = \frac{1}{2}$, so $\text{div}(t, E) = 1$, and the closer $\frac{1}{50} \sum_{(x, z) \in E_t} z$ is to 0 or 1, the smaller $\text{div}(t, E)$ is.

We base our similarity metric between search terms on Bing's lists of related searches. (We define diversity based on Google and similarity based on Bing so that our results are not concentrated on one search engine.) Whether the search engine suggests a given term depends on the similarity between that term and the original term. We assume that the higher in the list a term is, the more similar it is to the original search. Thus, we define the distance $\text{dist}(t, t', E)$ (where $\text{dist}(t, t', E) = 1 - \sigma(t, t', E)$) to be the position of t' in the list of related searches given the original search term t , divided by the list's length, or 1 if t' is not in the list.

In Figure 3, we display plots for two different searches. For the sake of readability, we only plot a subset of the related searches.

Figure 1: University admissions when each criterion is defined by an axis-aligned half-space. The orange bars indicate the number of white, non-Hispanic students whose features fall within a given range. The blue bars correspond to the black and Hispanic students. The dotted line in Figure 1b represents the original criterion (cumulative ACT score above 56). The solid line in Figure 1c represents the criterion returned by the system (class rank in the top decile).

We can see that some searches have nearby searches with diverse search results. For example, “neckerchief” is similar to “scarf”, but the images associated with “scarf” are almost all of women whereas the images associated with “neckerchief” have a more balanced gender ratio.

4.3 Job search

We next instantiate our general model for job applicant search. Given a recruiter’s original query such as “computer programmer”, the goal is to provide her with alternative search queries that return diverse sets of applicants.

To instantiate our general model, we define the set X of observable attributes to be an abstract set of job applicants’ observable attributes. For example, X could be a set of individuals’ LinkedIn profiles. Let $Z = \{0, 1\}$ be the set of sensitive attributes which indicate gender. In particular, given an example (x, z) , $z = 0$ if and only if x corresponds to a man. In this setting, we define the set of input examples E to be the entire set $X \times Z$. Each criterion $t \in T$ corresponds to an occupation. Given an individual (x, z) , we define $t(x)$ to be 1 if the individual is applying for the job t and 0 otherwise.

To measure diversity, we use workforce statistics from the [19]. Given an occupation t , these statistics tell us the fraction ω_t of full-time workers in the U.S. with that job year-round who are women. We define diversity to measure how far from $\frac{1}{2}$ that fraction is, i.e., $\text{div}(t, E) = 1 - |\frac{1}{2} - \omega_t|$. If there is a 50-50 gender split, then $\omega_t = \frac{1}{2}$ and $\text{div}(t, E) = 1$. Meanwhile, the closer ω_t is to 0 or 1, the smaller $\text{div}(t, E)$ is.

In order to characterize similarity among jobs, we use the Department of Labor’s Standard Occupational Classification (SOC) System, which organizes hundreds of occupations into a single hierarchy. Occupations are grouped together if they have similar duties or require similar skills, education, or training [20]. For a small subset of the hierarchy, see Figure 5 in the appendix. We define the similarity of two jobs based on the length of the path between them in the hierarchy. For two jobs t and t' , let $\ell(t, t')$ be the length of this path. The maximum path length between any two leaves in this hierarchy is 8. Thus, we define the substitutability of two jobs t and t' as $\sigma(t, t', E) = 1 - \frac{\ell(t, t')}{8}$ so that it is in $[0, 1]$.

In Figure 4, we display plots for two occupations. We can see that some occupations have nearby occupations with more diverse workforces. For example, “web developer” is similar to “computer programmer”, but there are more female web developers than computer programmers.

5 DIVERSITY ESTIMATION GUARANTEES

In this section, we provide provable guarantees for finding similar criteria with more diverse results. Here, our goal is to estimate the diversity of any given criterion using historical data. Calculating reliable diversity estimates is crucial because in some settings, it may be preferable or even mandatory to avoid using individuals’ protected attributes to measure diversity. For example, an admissions officer may not know the ethnicity of any college applicant from the current year. However, they may have access to historical data, such as the University of Texas data we use to run our experiments, which they can use to form these estimates. Furthermore, given estimates that reflect the diversity of any criterion over society as a whole, a user can introduce policies, such as a fixed college admissions criterion, that can stay in place for years at a time. For example, they can be confident that admitting students based on class rank will consistently yield more diverse student bodies than admitting students based on standardized test scores. Finally, using diversity estimates rather than true diversity scores in order to choose among criteria means that no individual will be favored based on his or her sensitive attributes.

Throughout this section, we consider a setting where a criterion t is a map $t : X \rightarrow \{0, 1\}$. Given an example (x, z) , $z \in [0, 1]$ is the degree to which x belongs to the minority group. Given a multi-set $E = \{(x_1, z_1), \dots, (x_N, z_N)\}$ of examples, we define $f_m(t, E)$ to be the degree to which an average instance $(x, z) \in E_t$ belongs to the minority group. In other words, $f_m(t, E) = \frac{1}{|E_t|} \sum_{(x, z) \in E_t} z$. Recall that in the university admissions setting, we define $\text{div}(t, E)$ to measure how far $f_m(t, E)$ is from a user-specified $c \in [0, 1]$, i.e., $\text{div}(t, E) = 1 - |c - f_m(t, E)|$. To justify this choice of $\text{div}(t, E)$, we prove that $f_m(t, E)$ estimates the protected attribute of an example

(a) Majority students' ACT scores versus class ranks.

(b) Minority students' ACT scores versus class ranks.

Figure 2: University admissions when each criterion is defined by the intersection of axis-aligned half-spaces. The plots display the majority (Figure 2a) and minority (Figure 2b) groups' ACT scores versus class ranks. The colorbars indicate the number of students who fall in each bin. The dotted lines represent the original criterion: ACT score above 56. The solid lines represent the suggested criterion: ACT score above 44 and class rank in the top decile.

$(x, z) \sim \mu$ conditioned on $t(x) = 1$. In other words, $f_m(t, E)$ approximates $\mathbb{E}_{(x,z) \sim \mu}[z \mid t(x) = 1]$ for all $t \in T$ so long as E has a sufficiently large number of candidates. That is, we show that with probability at least $1 - \delta$ over the draw of $E = \{(x_1, z_1), \dots, (x_N, z_N)\} \sim \mu^N$, for all $t \in T$ such that $\Pr_{(x,z)}[t(x) = 1] \geq c$, we have that $\left| f_m(t, E) - \mathbb{E}_{(x,z) \sim \mu}[z \mid t(x) = 1] \right| \leq \epsilon$ when N is sufficiently large with respect to ϵ , δ , and c .

We use the learning-theoretic notion of VC dimension [21] to provide this sample complexity guarantee. VC dimension measures the *intrinsic complexity* of binary-valued function classes, or in other words, classes of functions that map to $\{0, 1\}$. Given a set $\mathcal{G}$ of functions which map an abstract domain $\mathcal{A}$ to $\{0, 1\}$ and a distribution μ' over $\mathcal{A}$, bounding the VC dimension of $\mathcal{G}$ allows us to bound the number of samples $a_1, \dots, a_N \sim \mu'$ sufficient to ensure that the difference between the average value $\frac{1}{N} \sum_{i=1}^N g(a_i)$ of any

(a) Original search: “scarf”

(b) Original search: “haircut”

Figure 3: Plots for image search. The position on the x -axis of a image search t' equals its position in Bing's list of search terms related to the original search t . Its position on the y -axis equals the diversity $\text{div}(t, t', E)$. The line illustrates the Pareto frontier of this bi-criteria optimization problem.

function $g \in \mathcal{G}$ over the samples and its expected value $\mathbb{E}_{a \sim \mu'}[g(a)]$ is small. VC dimension applies specifically to binary-valued functions, so it does not immediately apply to our setting because the protected attributes z are real-valued and we are concerned with the conditional expectation $\mathbb{E}[z \mid t(x) = 1]$, not the expectation $\mathbb{E}[z]$. Nonetheless, we show how to use the VC dimension of T to bound the number of samples sufficient to ensure the difference between $f_m(t, E)$ and $\mathbb{E}[z \mid t(x) = 1]$ is small, for any criterion $t \in T$. In many applications, bounding the VC dimension of T is simple, so applying this sample complexity bound is straight-forward.

Below, we define VC dimension in terms of an abstract set of functions $\mathcal{G}$ which map a domain $\mathcal{A}$ to $\{0, 1\}$.

Definition 5.1 (VC dimension). We say that $\mathcal{G}$ shatters the set $\mathcal{S} = \{a_1, \dots, a_M\} \subseteq \mathcal{A}$ if for all binary vectors $\mathbf{b} \in \{0, 1\}^M$, there is a function $g_{\mathbf{b}} \in \mathcal{G}$ such that $g_{\mathbf{b}}(a_i) = b[i]$ for all $i \in [M]$. The VC dimension of $\mathcal{G}$, denoted $\text{VCdim}(\mathcal{G})$, is the size of the largest set that $\mathcal{G}$ shatters.

(a) Original search: “computer programmers”

(b) Original search: “social workers”

Figure 4: Plots for job search. The position on the x -axis of a job t' equals the path length between it and the original search t' in the SOC hierarchy. Its position on the y -axis equals the diversity $\text{div}(t, t', E)$. The line illustrates the Pareto frontier of this bi-criteria optimization problem.

For example, recall the college admissions criteria we study: axis-aligned half-spaces and the intersection of axis-aligned half-spaces in $\mathbb{R}^d$. For both, $\text{VCdim}(T) \in O(d)$.

We now present our sample complexity bound. The proof is in the full version of the paper.

THEOREM 5.2. *For any $\epsilon \in (0, 1/2)$ and $\delta \in (0, 1)$, with probability $1 - \delta$ over the draw $E = \{(x_1, z_1), \dots, (x_N, z_N)\} \sim \mu^N$, for all $t \in T$ such that $\Pr_{(x,z)}[t(x) = 1] \geq c$, $\left| f_m(t, E) - \mathbb{E}_{(x,z) \sim \mu}[z \mid t(x) = 1] \right| \leq \epsilon$ when $N \geq N_0 \in O\left(\frac{1}{c\epsilon^2} \left(\text{VCdim}(T) \log \frac{1}{\epsilon} + \log \frac{1}{\delta} \right)\right)$.*

5.1 Estimation with correlated examples

Theorem 5.2 applies when the examples (x, z) are identically and independently distributed. In this section, we provide diversity estimation guarantees even when the examples are correlated. Rather than the i.i.d. assumption, we assume there is a distribution $\bar{\mu}$ over datasets E of a fixed size m , so the support of $\bar{\mu}$ is a subset of $(X \times Z)^m$. We show how to use a set $\mathcal{S} = \{E^{(1)}, \dots, E^{(N)}\} \sim \bar{\mu}^N$ to estimate

$\mathbb{E}_{E \sim \bar{\mu}}[f_m(t, E)]$ for any criterion $t \in T$. One way to interpret this result is in the setting where we wish to design admissions criteria for a network of universities, such as all U.S. public community colleges. If we have a criterion that admits diverse students over a random set of N colleges, we can be confident the criterion also admits diverse students on average nationwide. Colleges may have correlated applicants depending on geography, and this section’s results allow for this correlation. Specifically, we prove the following guarantee. The proof is in the full version of the paper.

THEOREM 5.3. *For any $\epsilon, \delta \in (0, 1)$,*

$$N \geq N_0 \in \Theta\left(\frac{1}{\epsilon^2} \left(\text{VCdim}(T) \log(m \text{VCdim}(T)) + \log \frac{1}{\delta} \right)\right)$$

samples are sufficient to ensure that with probability $1 - \delta$ over the draw of $E^{(1)}, \dots, E^{(N)} \sim \bar{\mu}$, for all $t \in T$,

$$\left| \frac{1}{N} \sum_{i=1}^N f_m(t, E^{(i)}) - \mathbb{E}_{E \sim \bar{\mu}}[f_m(t, E)] \right| \leq \epsilon.$$

6 LIMITATIONS

We now describe several limitations of our work, some of which suggest directions for future research.

Scholars [10] have argued that *color-blind admissions criteria*, like those we study, are less effective at admitting students with strong academic performance compared to race-aware criteria. However, race-aware criteria may not be legal in many scenarios, in which case our approach presents a promising alternative.

A second limitation is that it may be computationally expensive to search for similar criteria yielding more diverse results. For example, this is the case in our university admissions application when each criterion is defined by the intersection of half-spaces. In fact, we prove that this problem is NP-complete in the appendix. Formulating an approximation algorithm is a promising direction for future work.

Lastly, if there is no clear notion of substitutability among criteria, then our framework will not apply. However, we believe our examples illustrate that there are often straight-forward ways to define substitutability.

7 CONCLUSION

We study the problem of suggesting similar criteria that yield more diverse results. In our setting, a domain expert chooses some selection criterion, such as an admissions criterion to choose among a set of college applicants. We do not assume she knows the true quality of any given criterion, or that the notion of “true quality” is well-defined. Rather, she uses heuristics and intuition to select a criterion she believes will return high-quality results. At the same time, this expert would like to ensure the results are diverse, but optimizing for diversity and quality may be difficult without access to ground-truth quality metrics. We introduce automated techniques for suggesting similar criteria yielding more diverse results. We demonstrate its strong performance in three critical domains: college admissions, image search, and job search. Finally, we complement our experiments with theoretical guarantees, analyzing the amount of historical data required to estimate criteria diversity.

REFERENCES

- [1] Alekh Agarwal, Alina Beygelzimer, Miroslav Dudík, John Langford, and Hanna Wallach. 2018. A reductions approach to fair classification. *Proceedings of the International Conference on Machine Learning (ICML)* (2018).
- [2] Asia J Biega, Krishna P Gummadi, and Gerhard Weikum. 2018. Equity of Attention: Amortizing Individual Fairness in Rankings. *Conference on Research and Development in Information Retrieval (SIGIR)* (2018).
- [3] Tolga Bolukbasi, Kai-Wei Chang, James Y. Zou, Venkatesh Saligrama, and Adam Tauman Kalai. 2016. Man is to Computer Programmer as Woman is to Homemaker? Debiasing Word Embeddings. In *Proceedings of the Annual Conference on Neural Information Processing Systems (NIPS)*.
- [4] Aylin Caliskan, Joanna J. Bryson, and Arvind Narayanan. 2017. Semantics derived automatically from language corpora contain human-like biases. *Science* 356, 6334 (2017), 183–186.
- [5] L Elisa Celis, Vijay Keswani, Damian Straszak, Amit Deshpande, Tarun Kathuria, and Nisheeth K Vishnoi. 2018. Fair and Diverse DPP-based Data Summarization. *arXiv preprint arXiv:1802.04023* (2018).
- [6] L. Elisa Celis, Damian Straszak, and Nisheeth K. Vishnoi. 2018. Ranking with Fairness Constraints. In *Proceedings of the International Colloquium on Automata, Languages and Programming (ICALP)*.
- [7] Flavio Chierichetti, Ravi Kumar, Silvio Lattanzi, and Sergei Vassilvitskii. 2017. Fair Clustering Through Fairlets. In *Proceedings of the Annual Conference on Neural Information Processing Systems (NIPS)*.
- [8] Cynthia Dwork, Moritz Hardt, Toniann Pitassi, Omer Reingold, and Richard Zemel. 2012. Fairness Through Awareness. In *Proceedings of the ACM Conference on Innovations in Theoretical Computer Science (ITCS)*.
- [9] Cynthia Dwork, Nicole Immorlica, Adam Tauman Kalai, and Mark DM Leiserson. 2018. Decoupled Classifiers for Group-Fair and Efficient Machine Learning. In *Proceedings of the Conference on Fairness, Accountability and Transparency (FAT*)*.
- [10] Roland G Fryer Jr, Glenn C Loury, and Tolga Yuret. 2007. An economic analysis of color-blind affirmative action. *The Journal of Law, Economics, & Organization* 24, 2 (2007), 319–355.
- [11] Moritz Hardt, Eric Price, and Nati Srebro. 2016. Equality of Opportunity in Supervised Learning. In *Proceedings of the Annual Conference on Neural Information Processing Systems (NIPS)*.
- [12] Rosie Jones, Benjamin Rey, Omid Madani, and Wiley Greiner. 2006. Generating query substitutions. In *Proceedings of the International World Wide Web Conference (WWW)*.
- [13] Chen Karako and Putra Manggala. 2018. Using Image Fairness Representations in Diversity-Based Re-ranking for Recommendations. In *Proceedings of the Workshop on Fairness, Accountability, and Transparency in Machine Learning (FATML)*.
- [14] Matthew Kay, Cynthia Matuszek, and Sean A Munson. 2015. Unequal representation and gender stereotypes in image search results for occupations. In *Proceedings of the Conference on Human Factors in Computing Systems (CHI)*.
- [15] Lyndon S Kennedy and Mor Naaman. 2008. Generating diverse and representative image search results for landmarks. In *Proceedings of the International World Wide Web Conference (WWW)*.
- [16] Jahna Otterbacher, Jo Bates, and Paul Clough. 2017. Competent men and warm women: Gender stereotypes and backlash in image search results. In *Proceedings of the Conference on Human Factors in Computing Systems (CHI)*.
- [17] Ashudeep Singh and Thorsten Joachims. 2018. Fairness of Exposure in Rankings. *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (2018).
- [18] Texas Higher Education Opportunity Project. 2008. Administrative College Application Data: Documentation for Public Use Data Files. https://theop.princeton.edu/admin_doc.html.
- [19] U.S. Census Bureau. 2016. Full-Time, Year-Round Workers and Median Earnings: 2000 and 2013-2016. <https://www.census.gov/data/tables/time-series/demo/industry-occupation/median-earnings.html>.
- [20] U.S. Office of Management and Budget. 2018. Standard Occupational Classification Manual. https://www.bls.gov/soc/2018/soc_2018_manual.pdf.
- [21] Vladimir Vapnik and Alexey Chervonenkis. 1971. On the uniform convergence of relative frequencies of events to their probabilities. *Theory of Probability and its Applications* 16, 2 (1971), 264–280.
- [22] Meng Wang, Kuiyuan Yang, Xian-Sheng Hua, and Hong-Jiang Zhang. 2010. Towards a relevant and diverse search of social images. *IEEE Transactions on Multimedia* 12, 8 (2010), 829–842.
- [23] Ke Yang and Julia Stoyanovich. 2017. Measuring fairness in ranked outputs. In *Proceedings of the Conference on Scientific and Statistical Database Management (SSDBM)*.
- [24] Ke Yang, Julia Stoyanovich, Abolfazl Asudeh, Bill Howe, HV Jagadish, and Gerome Miklau. 2018. A Nutritional Label for Rankings. In *Proceedings of the Conference on Management of Data (SIGMOD)*.
- [25] Muhammad Bilal Zafar, Isabel Valera, Manuel Gomez Rodriguez, and Krishna P Gummadi. 2017. Fairness constraints: Mechanisms for fair classification. *Proceedings of the International Conference on Artificial Intelligence and Statistics (AISTATS)* (2017).
- [26] Meike Zehlike, Francesco Bonchi, Carlos Castillo, Sara Hajian, Mohamed Megahed, and Ricardo Baeza-Yates. 2017. Fa*ir: A fair top-k ranking algorithm. In *Proceedings of the Conference on Information and Knowledge Management (CIKM)*.

A ADDITIONAL RELATED WORK

From a computational perspective, researchers studying algorithmic fairness have developed many algorithms assuming access to ground truth quality scores of individuals or outcomes. For example, in fair binary classification, a learning algorithm receives as input a training set of examples (x, y, z) where x is a set of observable attributes, z is a set of protected attributes, and y is a ground truth binary label. The goal is to learn a binary classifier which nearly matches the training set's ground truth labels, and which satisfies a given notion of fairness (e.g., [1, 9, 11, 25]). Thus, the quality of a classifier is clear: it measures how well that classifier matches the training set or a test set of fresh examples, potentially with the addition of a fairness loss function. In our setting however, the true quality of any given criterion may not even be defined.

Several papers have studied fair machine learning in the context of unsupervised learning, though in settings that are orthogonal to ours. These include works on fair clustering [7] and bias in word embeddings [3, 4].

B ETHICAL DATA USAGE

In this section, we discuss the ethics of our data usage. We use human data in two settings: college admissions and image search.

College admissions data. We use college admissions data collected by the Texas Higher Education Opportunity Project (THEOP), which is a part of Princeton University's Office of Population Research. We use a public use dataset which consists of college application data collected from the University of Texas at Austin, spanning from 1991 to 2003. Princeton's Institutional Review Board approved the administrative data collection and also the public release files. We requested access to these files and were granted it.

According to the data documentation [18], THEOP has taken steps to de-identify the data, including:

- Eliminating university-assigned applicant identification numbers and high school names, cities, and states.
- Eliminating small frequency (less than 20) cells by collapsing multiple values into range categories. For example, individual test score values are collapsed into test score ranges.
- Setting some values to missing to preserve the well-established categories but to hide individual values.

Nevertheless, it might be possible to re-identify individuals in this dataset. However, since this data is publicly available and (to the best of our knowledge) has been used by researchers in hundreds of academic works, we believe that our use of the dataset will not increase the risk of re-identification.

Image search. The other setting in which we use human data is image search. We do not download or store any images – we only store statistics in the form of Figure 3. We do not believe there is any risk of re-identification from these general statistics.

Figure 5: A small subset of the Department of Labor’s Standard Occupational Classification System.