



Mining the Minoria: Unknown, Under-represented, and Under-performing Minority Groups*

Mohsen Dehghankar
University of Illinois Chicago
mdehgh2@uic.edu

ABSTRACT

Due to a variety of reasons, such as privacy, data in the wild often misses the grouping information required for identifying minorities. On the other hand, it is known that machine learning models are only as good as the data they are trained on and, hence, may underperform for the under-represented minority groups. The missing grouping information presents a dilemma for responsible data scientists who find themselves in an unknown-known situation, where not only do they not have access to the grouping attributes but do not also know what groups to consider.

This paper is an attempt to address this dilemma. Specifically, we propose a minority mining problem, where we find vectors in the attribute space that reveal potential groups that are under-represented and under-performing. Technically speaking, we propose a geometric transformation of data into a dual space and use notions such as the arrangement of hyperplanes to design an efficient algorithm for the problem in lower dimensions. Generalizing our solution to the higher dimensions is cursed by dimensionality. Therefore, we propose a solution based on smart exploration of the search space for such cases. We conduct comprehensive experiments using real-world and synthetic datasets alongside the theoretical analysis. Our experiment results demonstrate the effectiveness of our proposed solutions in mining the unknown, under-represented, and under-performing minorities.

PVLDB Reference Format:

Mohsen Dehghankar and Abolfazl Asudeh. Mining the Minoria: Unknown, Under-represented, and Under-performing Minority Groups. PVLDB, 18(5): 1453 - 1465, 2025.

doi:10.14778/3718057.3718072

PVLDB Artifact Availability:

The source code, data, and/or other artifacts have been made available at https://github.com/UIC-InDeXLab/Mining_U3Ms.

1 INTRODUCTION

Specifying the scope of use of a dataset is an important step before sharing it [41]. Particularly, it is known that data-driven models are not reliable and may be unfair towards the groups that are not well represented in their (training) dataset [36, 37]. Therefore, it is important to identify the potential groups that are *under-represented*

*This work was supported in part by the National Science Foundation, Grant No. 2348919 and 2107290.

This work is licensed under the Creative Commons BY-NC-ND 4.0 International License. Visit <https://creativecommons.org/licenses/by-nc-nd/4.0/> to view a copy of this license. For any use beyond those covered by this license, obtain permission by emailing info@vldb.org. Copyright is held by the owner/author(s). Publication rights licensed to the VLDB Endowment.

Proceedings of the VLDB Endowment, Vol. 18, No. 5 ISSN 2150-8097.
doi:10.14778/3718057.3718072

Abolfazl Asudeh
University of Illinois Chicago
asudeh@uic.edu

and for which the models *under-perform*. However, *data in practice often misses to include demographic information of entities*. Even in cases where some demographic information are available, many others may be missing that is necessary for identifying representation bias [38]. To further motivate this problem, let us consider the following example:

EXAMPLE 1. A data-sharing platform (e.g., Chicago Open Data Portal [2]) would like to responsibly specify the under-represented and under-performing groups in their shared data in order to limit their scope of use. However, their datasets either do not include grouping attributes (such as race) or only contain some of those. Moreover, targeting a comprehensive audit, they do not want to limit their scope to a small set of predefined groups. Instead, they want to be proactive in their detection process, even for the groups they might have missed to consider.

For some cases such as image datasets, when a small and predetermined set of grouping attributes (e.g., gender) are missing, one may be able to use techniques such as crowdsourcing to identify under-represented groups based on those attributes [31]. Our problem, however, is that the grouping attributes are not predetermined. In other words, we consider an *unknown unknown* situation where not only do we lack access to the grouping attributes but we also do not know the groups of interest.

To address this problem, we introduce the problem of mining the potential groups that are under-represented and under-performing, when the grouping attributes are unknown. We refer to this problem as *Minoria mining*. To the best of our knowledge, this paper is *the first to introduce and study the Minoria mining problem*. In simple terms, the main question we want to answer in this paper is, “are there regions in data-space, representing potential groups that are (a) under-represented, for which (b) the model under-performs?” For instance, aligned with Example 1, our experiments on a Crimes dataset (§ 7.1.1) identified North-side Chicago as under-represented and under-performing. This information helps the model developers and data owners such as Chicago Open Data Portal to proactively identify a potential issue, decide if it is socially concerning, and limit the dataset’s scope-of-use accordingly.

Similar to any other data mining problem, multiple approaches may come to mind for addressing this problem. For example, one can adapt a clustering technique (such as k-means clustering [28] or subspace and projective clustering [24, 29, 32]) to cluster the data into potential groups. However, as we shall show in our experiments (§ 7.1.2) such an approach fails to identify under-represented groups, simply because there are not enough instances from those groups in the dataset. This is a well-known drawback of the clustering methods, that their performance degrades substantially for detecting small clusters with insufficient data points [34].

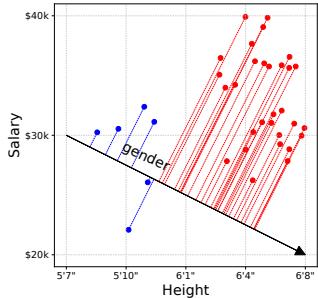


Figure 1: Toy example: Minoria mining using high-skew projections.

Instead, in this paper, we propose a novel approach that involves identifying high-skew projections of data along a specific direction in the feature space. Before outlining our technical contributions, we acknowledge that our approach is *only one of the possible solutions* for Minoria mining. There may be other complementary and equally effective techniques that address this problem. We hope that this paper will pave the way for future research to explore alternative approaches.

1.1 Technical Contributions

We propose a Minoria mining approach that assists its users to mine unknown, under-represented, and under-performing groups as *combinations of other features*.

Consider projecting data on a combination¹ of features, such that (i) the projected data is highly skewed (hence its tail indicating an under-represented group candidate) and (ii) the prediction loss at the tail is high. Our mining task is to discover such cases and *return the under-performing samples* (e.g., the misclassified samples) in the tail to the human expert for further investigation to verify if those represent a socially-valid group.

To better explain our approach, let us consider a toy example of a basketball dataset with several observation features X , including height and salary, and one target attribute (y) performance-score, where the goal is to predict y using X . Let the (height, salary) values be as shown in Fig. 1. Our system discovers the black vector as a combination of height and salary, where the projection of values on it are highly skewed. That is, after orthogonally projecting each point on the black vector, the distribution of the projected values is highly skewed. In this example, the projections are specified with dashed line, while the point in the head and tail of the distribution colored in red and blue, respectively. The points in the tail are considered as a *potential group* that are under-represented. Next, we observe that the prediction performance is significantly less for this potential group. We then report the misclassified samples on the tail as a potential Minoria to the user for further investigation. In this example, the discovered group represents (mostly) the female athletes that are not well-represented in the dataset.

Summary of contributions. In this paper, we introduce and formalize the Minoria mining problem. We propose an approach for the problem based on the high-skew projections in the feature space that reflect performance-disparities at the projection tail. It turns

¹We interchangeably use the term “linear direction” or simply “direction” for “linear combination”.

out that finding such projections is challenging as it requires identifying the median of each projection. To address this challenge, we utilize computational geometry notions such as duality, k -th level of arrangement, and the arrangement skeleton, and define the concept of median regions [21]. We prove the sufficient conditions for finding the maximum skew within each median region, which enables developing an efficient Ray-sweeping algorithm for a constant number of dimensions. Our Ray-sweeping algorithm has *the same time complexity* as of the k -th level of arrangement enumeration. This algorithm, however, suffers from the curse of dimensionality. Therefore, we propose practical solutions such as search-space discretization and smart exploration of the search space for large-scale high-dimensional settings. Besides theoretical guarantees, our experiments on multiple datasets confirm the validity of our proposal and the performance of our algorithms. For example, our experiments on a Chicago crime dataset discovered an under-performing high-skew projection that (mostly) reflects the white group.

2 PRELIMINARIES

Notations. We use $\{.\}^n$ to show a set of size n , and $[n]$ for the set of integers $\{1, \dots, n\}$. We also show the inner product of two vectors t and f as $t^\top f$.

Data Model. We consider a dataset $\mathcal{D} = \{t_i\}^n$ as a collection of n sample points (aka tuples). Each sample $t_i \in \mathcal{D}$ is defined as a tuple $t_i = \langle X, y \rangle$. $X = \langle x_1, \dots, x_d \rangle$, is a d -dimensional vector of Real values that corresponds with a set of *observation* attributes, also known as features. The observation attributes are used for prediction. For each sample $t_i \in \mathcal{D}$, we use $\langle t_{i,1}, \dots, t_{i,d} \rangle$, to show its features. The dataset also contains a *target attribute* (aka label) y .² The dataset \mathcal{D} is used to train a model $h_\theta(X)$ that predicts y using the observation attributes.

Group Membership. We consider a universe of (demographic) groups such as `male`, `female`, `Black`, `White`, etc., to which a subset of points belong to. For a group g , we use $\mathcal{D}^g = g \cap \mathcal{D}$ to specify the set of points that belong to it and $\mathcal{D}^{!g} = \mathcal{D} \setminus \mathcal{D}^g$ as the points that do not belong to it. A grouping (aka, sensitive) attribute separates the group g from the others. For example, the grouping attribute `gender` separates the `female` individuals (\mathcal{D}^f) from others ($\mathcal{D}^{!f}$). We assume *the dataset does not include the grouping attributes*. Moreover, *the set of grouping attributes of interest is not predetermined*. In other words, not only do we not know the grouping information, we do not also know what groups we should consider.

2.1 Problem Definition

Minoria Mining. Our objective is to mine the unknown, under-represented, and underperforming minority groups (referred to as Minoria). Specifically, given a dataset \mathcal{D} , identify groupings of \mathcal{D} into \mathcal{D}^g and $\mathcal{D}^{!g}$ such that the following conditions hold³:

- (a) g is under-represented in the dataset, i.e., $|\mathcal{D}^g| \ll |\mathcal{D}|$.

²The extension of our solutions for datasets with more than one target attribute is straightforward by solving the problem separately for each attribute y_k .

³Note that, slightly abusing the notation, we also use g for the group candidates. A discovered group candidate may or may not express a socially valid group.

(b) Predictions based on \mathcal{D} are not accurate for g . Formally, let h_θ be the optimal model that minimizes the expected loss $\mathbb{E}[L_{\mathcal{D}}(\theta)]$ on \mathcal{D} . Then, for a given performance-disparity threshold τ , $\mathbb{E}[L_{\mathcal{D}^g}(\theta)] - \mathbb{E}[L_{\mathcal{D}}(\theta)] \geq \tau$.

We combine X to form *proxy* attributes that identify Minoria. Formally, we would like to find projections of points in \mathcal{D} from \mathbb{R}^d to \mathbb{R} such that the projection (almost) separates \mathcal{D}^g from \mathcal{D}^{g^c} . Furthermore, to satisfy condition (a), we are interested in finding projections with *high skew* on \mathcal{D} , indicating an under-represented group candidate at the tail of the projection.

Having identified such a group candidate g , we return the under-performing samples in g (the tail of a high-skew projection) to the human expert for further investigation.

Without loss of generality, we focus on linear projections in this paper and use *Pearson's median skewness* [19] for computing the skew on a projection. The Pearson's median skewness of a set of values $V = \{v_i\}^n$, where $v_i \in \mathbb{R}$ is computed as:

$$\text{skew}(V) = \frac{3(\mu(V) - \nu(V))}{\sigma(V)} \quad (1)$$

Where $\mu(V)$, $\nu(V)$, and $\sigma(V)$ are mean, median, and standard deviation of V . For a unit vector f , we define \mathcal{D}_f as the projection of all points on f :

$$\mathcal{D}_f = \{t_i^\top f \mid t_i \in \mathcal{D}\} \quad (2)$$

Then, the *unit* vector f^* with the highest skew is defined as $f^* = \arg \max_{f^\top f=1, f \in \mathbb{R}^d} \text{skew}(\mathcal{D}_f)$.

Since we would like to identify the Minoria candidates for further investigation, instead of returning the global optimum, i.e., the max-skew projection, we are interested in finding a top- ℓ high-skew optima. Putting everything together, we define the Minoria mining problem as follows:

PROBLEM 1 (MINORIA MINING). *Given a dataset \mathcal{D} of n samples, each containing a d -dimensional observation vector X and a target attribute y , and the performance-disparity threshold τ , find the unit vectors $\{f_1, \dots, f_\ell\}$ as*

$$\{f_1, \dots, f_\ell\} = \underset{f^\top f=1, f \in \mathbb{R}^d}{\text{top-}\ell \text{ skew}(\mathcal{D}_f)}$$

such that for all $i \in [\ell]$, $\mathbb{E}[L_{\mathcal{D}^{g_i}}(\theta)] - \mathbb{E}[L_{\mathcal{D}}(\theta)] \geq \tau$.
 top- ℓ refers to the top ℓ local optima (for a small constant ℓ) with the maximum skew, θ is the parameter of the model h_θ that minimizes the expected loss on \mathcal{D} , and g_i is the p -percentile tail of \mathcal{D}_{f_i} . \square

Addressing Problem 1 is challenging as it requires identifying the median of each possible projection while there are infinite possible projections defined using different vectors f . To tackle this challenge, we provide a geometric dual transformation of \mathcal{D} in the next section that helps us design our efficient algorithm in § 4.

3 GEOMETRIC INTERPRETATION

In this section, we present the geometric interpretation of data that forms the foundation of our solutions in later sections. First, in Section 3.1, we provide a notion of duality [21], where samples are presented as hyperplanes in a geometric space and prove some valuable characteristics of this transformation. We then consider the arrangement [21] of dual hyperplanes to identify high-skew projections efficiently.

3.1 Dual-Space Transformation

It is popular to present data in *primal space*, where each sample is shown as a point in \mathbb{R}^d . Instead, we use the following *dual space* transformation [9, 13] that the data points as a set of hyperplanes.

The dual space. Let $t_i = \langle t_{i_1}, t_{i_2}, \dots, t_{i_d} \rangle$ be a tuple in the dataset \mathcal{D} . We define the dual representation of t_i , as the hyperplane $d(t_i)$ in \mathbb{R}^d ($\{x_1, \dots, x_d\}$ are the d axes of the dual space):

$$d(t_i) : t_{i_1}x_1 + t_{i_2}x_2 + \dots + t_{i_d}x_d = 1 \quad (3)$$

As a running example in this section, consider a toy dataset $\mathcal{D} = \{t_1 \langle .5, 1.5 \rangle, t_2 \langle 1, .75 \rangle, t_3 \langle 2, 1 \rangle\}$. Fig. 2a shows \mathcal{D} as points in the primal space. In Fig. 2b, every t_i is represented as a line in the dual space. For instance, the point $t_1 \langle .5, 1.5 \rangle$ is shown as the line $d(t_1) : 0.5x_1 + 1.5x_2 = 1$.

Similarly, let H be a hyperplane in the primal space, specified with a set of d points p_1, \dots, p_d . The dual presentation of H is the point $d(H)$, specified as the intersection of $d(p_1), \dots, d(p_d)$. In the following, we review some properties of this dual space transformation and its relationship with the primal space. For ease of explanation, for every unit vector f , we define \vec{r}_f as the origin-anchored ray that passes through the point f .

PROPERTY 1. *Consider a point $t \in \mathbb{R}^d$ and a hyperplane H in the primal space. t lies on the hyperplane H if and only if the point $d(H)$ lies on the hyperplane $d(t)$.*

Let \vec{r}_f be the origin-anchored ray passing through $d(H)$ in the dual space. Also, let h be the intersection of H with the origin-anchored vector that is orthogonal to it in the primal space. Then, H is defined as the set of points $\{x : x^\top h = |h|\}$. As a result, $d(H) = h/|h|$. This holds because $H = \{x : d(H)^\top x = 1\}$. Then, the unit vector $f = \frac{h}{|h|}$, defining the ray \vec{r}_f , is the unit vector of H .

PROPERTY 2. *Let t be a point and f be a unit vector in \mathbb{R}^d . The intersection of \vec{r}_f and $d(t)$ is $f/(t^\top f)$.*

Property 2 holds because the intersection of $d(t)$ and \vec{r}_f can be shown as αf for a scalar value $\alpha \in \mathbb{R}$. Let αf intersect with the hyperplane $d(t)$, so αf satisfies Equation 3. Then,

$$t^\top (\alpha f) = 1 \Leftrightarrow \alpha t^\top f = 1 \Leftrightarrow \alpha = \frac{1}{t^\top f} \quad (4)$$

As a result of Property 2, for a pair of points t_i and t_j , $t_i^\top f \leq t_j^\top f$ if and only if the intersection of $d(t_i)$ lies further than $d(t_j)$ on the ray \vec{r}_f . This observation will result in the following Lemma.

LEMMA 1. *The ordering of points in \mathcal{D}_f is the same as the reverse order of intersections of $d(t_i)$ hyperplanes with \vec{r}_f . Hence, the dual-space transformation preserves the order of projected points.*

An illustration of these observations is shown in Fig. 2a and Fig. 2b for the direction $f = \langle \frac{\sqrt{2}}{2}, \frac{\sqrt{2}}{2} \rangle$. As shown in Fig. 2a, the projected points on f are ordered as $\{t_2, t_1, t_3\}$. One can confirm that the reverse ordering of the dual-line intersections with \vec{r}_f in Fig. 2b is also the same.

3.2 Median Region

A key concept that helps us to solve Problem 1 efficiently is the computational geometry notion of *k-th level of arrangement* [21]. Consider the dual hyperplanes $\mathcal{H} = \{d(t_1), \dots, d(t_n)\}$. The dissection of the dual space with the hyperplanes is called the *arrangement*

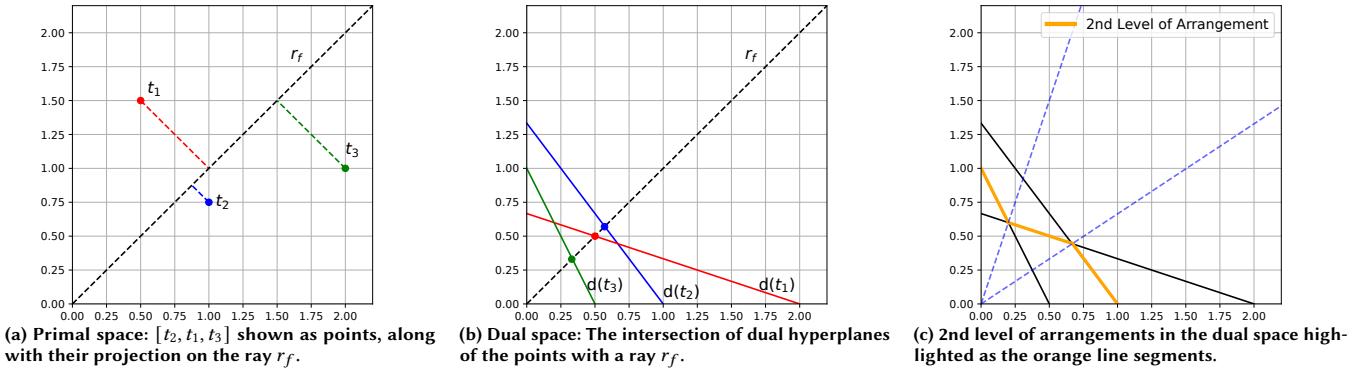


Figure 2: The illustration of the toy dataset $\mathcal{D} = \{t_1(.5, 1.5), t_2(1, .75), t_3(2, 1)\}$ in the primal space, the dual space, and the 2nd level of the arrangement in the first quadrant. The order of projection in Fig. 2a is the reverse of the intersection of dual hyperplanes with r_f ($[d(t_3), d(t_1), d(t_2)]$) in Fig. 2b. In Fig. 2c, the dotted blue lines indicate the boundaries of the median regions (the change in the line segment indicates a change in the median (2nd point) of \mathcal{D}_f).

\mathcal{A} of \mathcal{H} , which we simply call the arrangement. The intersection of $\delta \in [d]$ hyperplanes in \mathcal{H} forms a $(d - \delta)$ -face of the arrangement. In particular, the intersection of d hyperplanes is a (0) -face (a point), referred to as a *vertex* of the arrangement, while the (1) -faces (line segments) connecting two vertices are called the *edges*. We call two vertices that share an edge as *neighbors*. Please refer to [21] for more details about the arrangement.

Now consider the origin-anchored ray \vec{r}_f in the dual space and the intersections of the dual hyperplanes with it. Based on Lemma 1, the k -th closest intersection to the origin specifies the tuple with the k -th largest value in \mathcal{D}_f . We define the universe of all such intersection points for all vectors f as the k -th level of arrangement, referred as \mathcal{A}_k . In \mathbb{R}^2 , we show \mathcal{A}_k as the ordered list of the vertices (the intersection points) of the k -th level, in counter-clock-wise (from x -axis to y) or clock-wise.

For example, Fig. 2c shows (in orange) the 2nd level of an arrangement of 3 (toy) tuples in the dual space in \mathbb{R}^2 . In this example, $\mathcal{A}_2 = \langle 0, \times_{1,2}, \times_{1,3}, \frac{\pi}{2} \rangle$, where $\times_{i,j}$ is the intersection between $d(t_i)$ and $d(t_j)$. One can see that the k -th level is a chain of line segments forming a non-convex shape. In addition, the set of top- k tuples in each line segment is fixed and is different from the neighboring line segments. Moreover, when $k = \frac{n}{2}$, the k -th level of the arrangement tracks the median for any possible direction f . Specifically, the intersection of $\mathcal{A}_{\frac{n}{2}}$ with a ray \vec{r}_f is a line segment, specifying the median of \mathcal{D}_f . Evidently, the median of all rays intersecting with the same line segment in $\mathcal{A}_{\frac{n}{2}}$ is the same. Following this observation, we define the notion of median regions as follows.

DEFINITION 1 (MEDIAN REGION). A median region is the range of rays \vec{r}_f , for which the median of the projection \mathcal{D}_f is the same.

The partitioning of the origin-anchored rays in the dual space according to the $(d - 1)$ -dimensional simplexes (line-segments in 2D) of $\mathcal{A}_{\frac{n}{2}}$ provides the set of all median regions. Specifically, in 2D each median region is a range that begins/ends at the beginning/end of each line segment (e.g., see Fig. 2c). Moreover, since within every median region, no dual hyperplane intersects with the median simplex, the set of tuples with projection values being larger (resp. smaller) than the median does not change. For example, in Fig. 2c,

consider the median region between the x -axis and the first (bottom) blue ray. $d(t_2)$ is the median line-segment. The intersection of $d(t_3)$ with any ray in this region remains below the median (closer to the origin). Similarly, the intersection of $d(t_1)$ above the median.

Using these observations, one can write a linear program for every median region to find the maximum skew within a median region. Instead, we introduce Theorem 2 as a key observation that enables a more efficient algorithm in Section 4. Let $\mu(\mathcal{D})$ be the vector in \mathbb{R}^d that contains the mean of \mathcal{D} , i.e., $\mu(\mathcal{D}) = \frac{1}{n} \sum_{i=1}^n t_i$. We define $q_i = t_i - \mu(\mathcal{D})$. In other words, q_i 's are the mean-normalized tuples of \mathcal{D} . Let Q be a $d \times n$ matrix where each column of Q is q_i^\top . For a unit vector f , let t_{mf}^\top be the median of \mathcal{D} after projection on f . In other words, $t_{mf}^\top f$ is the median of \mathcal{D}_f . Theorem 2 helps us to find the high skew direction vector f^* efficiently⁴:

THEOREM 2. For a given median region with median t_{mf} , the max skew direction f^* either falls in the boundaries of this median region $\frac{(QQ^\top)^{-1}q_{mf}}{\|(QQ^\top)^{-1}q_{mf}\|}$ or it is equal to $\frac{(QQ^\top)^{-1}q_{mf}}{\|(QQ^\top)^{-1}q_{mf}\|}$.

4 MINORIA MINING IN 2D

In this section, we start by studying the problem in \mathbb{R}^2 – referred to as the 2D case. Before providing the technical details, we would like to clarify that in the 2D case, \mathcal{D} can have more than two observation attributes (used for prediction). However, the projection directions are considered with respect to a pair of observation attributes called attributes of interest. For example, in our experiments in § 7, we use longitude and latitude as the attributes of interest for Minorria mining in a crime dataset.

Recall that our goal is to find a set of high-skew directions. Let $\mathcal{H} = \{d(t_i) | t_i \in \mathcal{D}\}$ be the set of dual hyper-planes in the dual space and let \mathcal{E} be an oracle that returns $\mathcal{A}_{\frac{n}{2}}$ (the ordered list of vertices in the $\frac{n}{2}$ -th level of arrangement of \mathcal{H}) in T_{enum} time.

Given $\mathcal{A}_{\frac{n}{2}}$, we can iterate through the neighboring median regions by iterating through the elements (vertices) of this ordered list. Based on Theorem 2, within each median region, the high-skew direction is either at the boundary points (two vertices of the line

⁴Due to the space limitations, the proof is provided in the technical report [16].

segment) or at a specific direction, which can be calculated in $O(d)$. Since this time is dominated by the time taken to calculate the skew value of the directions corresponding to the boundary points, we can ignore it while analyzing the complexity of our algorithm.

In the second phase of our algorithm, we should iterate through the neighboring median regions and calculate the skew value of the directions of the boundaries of these regions. We want to return the top- ℓ high-skew directions as the Minoria candidates. In order to do that, we use a Max-Heap data structure to push these directions based on their skew value. We define T_{skew} as the running time of this phase.

In the end, we pop the top- ℓ directions from the heap, evaluate the model on the tail of these directions, and check the error on the tail to see if it represents an under-represented group candidate. A visualization of this approach is provided in Figure 3.

Let T_{pop} be the time complexity of the final step. Then, the overall time complexity of the algorithm is $T = T_{enum} + T_{skew} + T_{pop}$.

Having discussed the overall description of the 2D solution, let us next explain the oracle \mathcal{E} for enumerating $\mathcal{A}_{\frac{n}{2}}$.

4.1 Enumerating the k -th Level

Our efficient algorithm (§ 4.3) takes $\mathcal{A}_{\frac{n}{2}}$, the $\frac{n}{2}$ -th level in an ordered list, as the input. The k -th level of arrangement in \mathbb{R}^2 has a complexity (the number of its vertices) of $m = O(n \cdot k^{\frac{1}{3}})$ [17]. There are efficient deterministic and randomized algorithms for enumerating the k -th level. The SOTA deterministic algorithm is by Timothy Chan [10]. It enumerates \mathcal{A}_k as an ordered list in $O(n \log m + m \log^{1+\alpha} n)$, where α is a small positive constant, n is the number of hyperplanes, and m is the number of vertices in the k -th level of arrangement. The randomized incremental algorithm performs this traversal in $O(m + n \log n)$ [11].

4.2 The Warm-up Algorithm

So far, we outlined the key steps in our algorithmic framework. Following these steps, the Warm-up algorithm⁵ first enumerates $\mathcal{A}_{\frac{n}{2}}$. It then computes the median by evaluating all projections along a specified direction. Finally, it identifies and retrieves the top skewed elements.

Complexity Analysis. The Warm-up algorithm conducts the three aforementioned steps separately. In the first step, it transforms the dataset into the dual hyperplanes and calls the oracle \mathcal{E} to enumerate $\mathcal{A}_{\frac{n}{2}}$. As a result, the first step takes $T_{enum} = O(n \log(nk^{\frac{1}{3}}) + nk^{\frac{1}{3}} \log^{1+\alpha} n) = O(n^{\frac{4}{3}} \log^{1+\alpha} n)$ for the deterministic algorithm, or $T_{enum} = O(nk^{\frac{1}{3}} + n \log n) = O(n^{\frac{4}{3}})$ when using the randomized incremental algorithm (note that $k = \frac{n}{2}$ here). In the second step, for each vertex in $\mathcal{A}_{\frac{n}{2}}$, it calculates the skew value and adds it to a max-heap. Calculating the skew requires calculating the mean, median, and standard deviation of points projected on the direction f and takes $O(n)$. Inserting into a max-heap using a Fibonacci heap takes a constant time. As a result, $T_{skew} = O(n^2 k^{\frac{1}{3}}) = O(n^{\frac{7}{3}})$. In the third step, the Warm-up algorithm pops the high skew directions from the heap and calculates the model's loss on the tail of each direction. Popping from a max-heap takes an amortized time of $O(\log n)$. Calculating the model accuracy would take $O(n)$. Hence, assuming that the number of

directions evaluated before returning the top- ℓ is a small constant, $T_{pop} = O(n \log n)$.

As a result, the time complexity of 2DWarm-up is dominated by T_{skew} and is equal to $T = T_{enum} + T_{skew} + T_{pop} = O(n^{\frac{7}{3}})$.

4.3 The Ray-sweeping Algorithm

The Warm-up algorithm is super-quadratic and not efficient. Therefore, in this section, we propose the Ray-sweeping algorithm that has *the same time complexity* as of the enumeration of $\mathcal{A}_{\frac{n}{2}}$. Similar to the Warm-up algorithm, the Ray-sweeping algorithm also follows the three steps of (1) enumerating $\mathcal{A}_{\frac{n}{2}}$, followed by (2) computing the medians and (3) identifying the high-skew projections. However, unlike Warm-up, it leverages our earlier observation to compute medians more efficiently by sweeping a ray counterclockwise from the x-axis to the y-axis. At a high level, it maintains some aggregates while sweeping over $\mathcal{A}_{\frac{n}{2}}$, which enables updating the skew values in constant time.

The key observation is that while moving between two neighboring median regions, *the set of dual lines above/below the median changes by only one element*. As a result, using the pre-calculated values for the previous vertex, we can calculate the skew value of the current vertex in constant time (except for the very first vertex in the list). Please Refer to § 5.2 for more details about the constant-time skew-value update process. Moreover, when visiting a new vertex in $\mathcal{A}_{\frac{n}{2}}$, the new median can be found in constant time simply because each vertex is the intersection of the medians of the two neighboring median regions. The pseudo-code of this approach is provided in Algorithm 1. A visualization of this algorithm on a toy example is provided in Figure 3.

Complexity Analysis. Steps one and three of the Ray-sweeping algorithm are similar to the Warm-up Algorithm. For the second part, we calculate each skew value in a constant time. As a result, the running time of the second step is $T_{skew} = O(nk^{\frac{1}{3}}) = O(n^{\frac{4}{3}})$. Consequently, the time complexity of the Ray-sweeping algorithm is no longer dominated by step two. Therefore, the time complexity of Ray-sweeping is equal to the time taken for enumerating the $\frac{n}{2}$ -th level of the arrangement – i.e., $O(n^{\frac{4}{3}} \log^{1+\alpha} n)$ and $O(n^{\frac{4}{3}})$ for the deterministic and randomized algorithms, respectively.

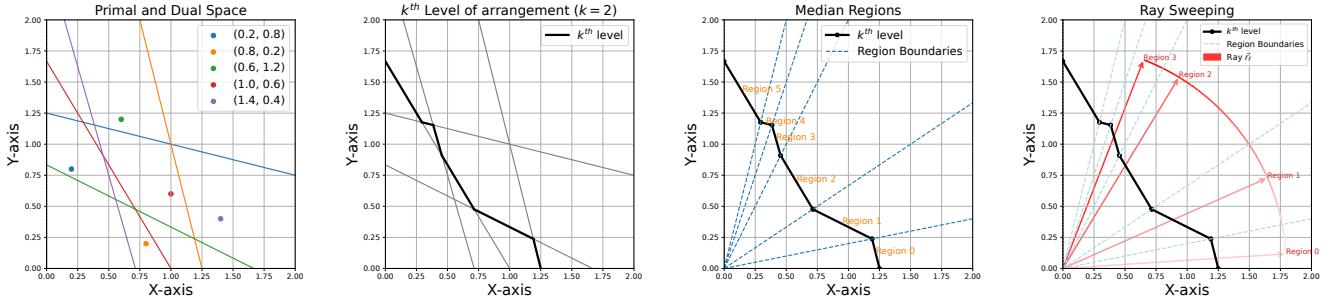
5 GENERALIZED RAY-SWEEPING FOR HIGHER DIMENSIONS

In this section, we extend our Ray-sweeping algorithm to the general setting, where $d > 2$ is a small constant.

5.1 Generalized Ray-sweeping Algorithm

For ease of explanation, we explain the Ray-sweeping algorithm in the first quadrant of \mathbb{R}^d , where all values are positive. When $d = 3$, the k -th level of the arrangement of n hyperplanes has a complexity of the $O(n \cdot k^{\frac{3}{2}})$ [39], while its complexity is bounded by $O(n^{\lfloor \frac{d}{2} \rfloor} k^{\lceil \frac{d}{2} \rceil})$ in general settings [5]. As a result, setting $k = \frac{n}{2}$, $\mathcal{A}_{\frac{n}{2}}$ has $O(n^{2.5})$ vertices when $d = 3$ and at most $O(n^d)$ vertices (each being is the intersection of $d - 1$ hyperplanes) when $d > 3$. The SOTA algorithm for enumerating the vertices of the k -th level of the arrangement while following the neighboring vertices is a randomized algorithm proposed by Agarwal et al. [4], with time

⁵Please find the pseudo-code of the Warm-up algorithm in the technical report [16].



(a) An example of 5 points and their dual lines in the same plane. (b) The k -th level of arrangement for $k = \frac{n}{2}$, representing the median points. (c) The median regions and their boundaries for a ray starting from the origin. (d) Starting from 5 points in the plane, the algorithm finds the $\frac{n}{2}$ -th level of arrangements for the dual lines. (e) Each median region is a line segment in this arrangement. The line segment represents the median point for any direction in this region. (f) The algorithm sweeps a ray from the x -axis to the y -axis for each region, calculates the skew value in this region efficiently, and pushes this value into a heap. In the end, the best directions, in terms of their skew, are extracted from the heap.

Algorithm 1 Generalized Ray Sweeping

```

1: procedure GENERALIZEDRAYSWEEPING( $\mathcal{D}, \theta$ )
2:    $\mathcal{H} \leftarrow \text{Dual}(\mathcal{D})$ ;  $\text{Output} \leftarrow []$ 
3:    $\mathcal{G}_{\frac{n}{2}} = \mathcal{E}(\mathcal{H})$   $\triangleright$  The graph of the  $k$ -level skeleton (vertices and edges).
4:    $\text{Heap} \leftarrow []$   $\triangleright$  Initialize a max heap.
5:   for  $v \in V(\mathcal{G}_{\frac{n}{2}})$  in a DFS traversal order do
6:      $f \leftarrow v/||v||$ 
7:      $\text{skew} = \text{GetSkew}(\mathcal{D}, f)$   $\triangleright$  Use the aggregated values to compute the skew in constant time.
8:      $\text{Heap.push}(\langle f, \text{skew} \rangle)$ 
9:   while  $\text{Output}$  size is less than  $l$  do
10:     $f \leftarrow \text{Heap.pop}()$ 
11:     $T \leftarrow p\text{-tail}(f)$ 
12:    if  $L_T(\theta) - L_D(\theta) > \tau$  then  $\text{Output.append}(f)$ 
13: return  $\text{Output}$ 

```

complexity of $O(n^{\lfloor \frac{d}{2} \rfloor} k^{\lceil \frac{d}{2} \rceil})$. Therefore, setting $k = \frac{n}{2}$, enumerating $\mathcal{A}_{\frac{n}{2}}$ in general setting takes $O(n^d)$.

After enumerating $\mathcal{A}_{\frac{n}{2}}$, we can follow the same aggregation process as the 2D space and calculate the skew value of each neighboring vertex in time $O(d)$. As a result, $T_{\text{skew}} = O(dn^d)$. Hence, given that $T_{\text{enum}} = O(n^d)$, the time complexity of the Ray-sweeping algorithm in the general setting is bounded by $O(dn^d)$.

The pseudo-code of the Ray-sweeping algorithm is provided in Algorithm 1. The algorithm first constructs the graph of k -th level of arrangements, aka the k -th level skeleton of the arrangement. In the skeleton graph, the vertices of \mathcal{A}_k are the nodes, while the edges are between the neighboring vertices. Please refer to Chapter 9 of [21] for more details about the skeleton of an arrangement and its Eulerian tour. After constructing the skeleton, the Ray-sweeping algorithm visits its vertices in a depth-first-search manner in $O(n^d)$.

5.2 Constant-time skew-value update process

In this section, we propose an aggregation method to update the skew of a direction corresponding to a vertex in the k -th level of arrangements in constant time with respect to n . Remember the Ray-sweeping algorithm explained in Section 4.3. We make a pass over $\mathcal{A}_{\frac{n}{2}}$, which is an ordered list of neighboring vertices in the k -th level of arrangement.

Let $[v_1, v_2, \dots, v_m]$ be this ordered list of vertices we visit one after the other. In other words, at each iteration $i \in [m]$, we explore the direction $f_i = \frac{v_i}{||v_i||}$. Refer to Algorithm 1 for the details. Let $\text{skew}_i = \text{skew}(f_i) = \frac{3(\mu_i - v_i)}{\sigma_i}$, where μ_i , v_i , and σ_i are the mean, median, and the standard deviation in the iteration i . The mean μ_i can be calculated in $O(d)$, because $\mu_i = \mu(\mathcal{D})^\top f_i$, where $\mu(\mathcal{D})$ is the mean point of the dataset. We need to calculate and store $\mu(\mathcal{D})$ in the pre-processing step. As a result, in the iteration i , we can calculate μ_i in constant time with respect to n (size of \mathcal{D}).

The median v_i can be calculated using the vertex v_i . Each vertex in the k -th level of arrangement is the intersection of multiple adjacent median regions. Each median region corresponds to a segment of a hyperplane, which is the dual transform of the median point for that region. Assume for each vertex in $\mathcal{A}_{\frac{n}{2}}$, we also store the set of median regions intersecting in that point (we use a hashmap to calculate this set of hyper-planes in constant time). This additional information is provided by the algorithms we discussed for constructions of the k -th level of arrangement. Now, given a hyperplane $d(t)$ in the k -th level of arrangement intersecting in v_i , we can calculate the median by using the inverse dual transform $d^{-1}(d(t))$ in $O(d)$ time.

$$\text{Now let us consider } \sigma_i: \quad \sigma_i = \sqrt{\frac{\sum_{j \leq n} (t_j^\top f_i - \mu_i)^2}{n}} \quad (5)$$

We need to calculate the sum $\sum_j (t_j^\top f_i - \mu_i)^2$ in constant time. We can rewrite this as:

$$\sum_j (t_j^\top f_i - \mu_i)^2 = \sum_j (t_j^\top f_i)^2 + 2\mu_i \sum_j t_j^\top f_i + n\mu_i^2 \quad (6)$$

As discussed before, the last term $n\mu_i^2$ can be calculated in constant time. For the second term, we need to store $S = \sum_j t_j$ in the pre-processing step. We can then calculate $2\mu_i S^\top f_i$ in constant time. For the first term, we define T as a $n \times d$ matrix having each t_j as its j -th row. Then we can rewrite this as:

$$\sum_j (t_j^\top f_i)^2 = f_i^\top T^\top T f_i \quad (7)$$

Where $T' = T^\top T$ is a $d \times d$ matrix, which can be calculated in the pre-processing step and stored. Then the matrix multiplication $f^\top T' f$ can be calculated in $O(d)$. As a result, skew_i can be calculated in $O(d)$ – constant time with respect to n .

6 PRACTICAL HEURISTICS

The generalized Ray-sweeping algorithm suffers from the *curse of dimensionality* [21]. Therefore, in this section, we propose practical solutions based on search-space discretization (§ 6.1) and smart exploration of the projection space (§ 6.2 and 6.3).

6.1 MD Discretization

Search-space discretization is common for addressing the curse of dimensionality. Following this idea, we propose two approaches to discretize the space of all possible directions f in \mathbb{R}^d and find the high-skew projections in the discretized space. Our first strategy is *uniform* discretization, where the continuous space of all directions f is partitioned into equal-size grids, and uniformly distributed samples are taken from the grid. The second strategy works by incrementally building a set of l directions using *Quadratic Programming* (QP), where the constraints of the QP are adjusted to enforce exploring the directions farthest from each other.

6.1.1 The Grid-partitioning Approach. Like the previous sections, we only focus on the positive partition of \mathbb{R}^d . Using the Polar coordinate system, we can uniquely specify each unit vector f in this space by $d - 1$ angles $\{\alpha_1, \alpha_2, \dots, \alpha_{d-1}\}$ between f and the positive direction of x_i axes ($i \in [d - 1]$). Here $\alpha_i \in [0, \frac{\pi}{2}]$. In other words, f comes from an $(d - 1)$ -dimensional space spanned by $(d - 1)$ independent angles. Now, we can grid-partitioning this space by dividing each angle into equi-size ϵ -length intervals. As a result, the total number of grid-cells would be $(\frac{\pi}{2} \cdot \frac{1}{\epsilon})^{d-1}$.

According to the problem definition, we want to report top- l high-skew directions f as candidates to find Minoria. We should sample at least l directions from the grids to achieve this. We want our samples to span all possible directions, so we should follow a sampling technique to satisfy the diversity of directions.

We follow the DPP (Determinental Point Processing) sampling schema to achieve this [6]. Using this approach, one can get a set of l diverse samples from a universe of n points in $O(n \cdot l^{\omega-1})$, where $\omega \approx 2.37$ is the matrix multiplication factor. In this approach, we sample $O(l)$ grids in time of $O((\frac{\pi}{2} \cdot \frac{1}{\epsilon})^{d-1} \cdot l^{\omega-1})$ and push them in a Max-Heap structure based on the skew value. We analyze each grid by popping the heap and filtering those satisfying the condition of Problem 1. If we find less than l such directions, we continue sampling more directions (grids).⁶ Calculating the skew of l directions takes an additional $O(l \cdot n \cdot d)$ time.

6.1.2 The Candidate-group diversification Approach. In this approach, we build a diverse set of candidate directions incrementally until we reach the l direction that satisfies the problem requirements. We start from an initial set of unit directions $P_0 = \{u_1, u_2, \dots, u_d\}$ where u_j is a unit vector lying on the j -th axis (x_j).

At each iteration i , we construct the vector set P_i by adding a new vector $p^{(i)}$ to P_{i-1} . This vector $p^{(i)}$ is as far as possible from the other vectors (alternatively, each vector can also be considered as a point with a unit distance from the origin). We can find such a point by solving the following *Quadratic Program*:

$$\begin{aligned} \min y \\ \text{s.t. } p_j^\top p^{(i)} \leq y \quad \forall j : p_j \in P_{i-1} \\ |p^{(i)}| = 1 \end{aligned} \quad (8)$$

⁶The bottleneck of this algorithm is the sampling step.

After solving the above QP, we will calculate the skew value for the direction $p^{(i)}$ in $O(dn)$ time. Solving a quadratic program with n variables has a time complexity of $O(n^{3.67} \cdot \log n \cdot \frac{1}{\epsilon} \cdot \log \frac{1}{\epsilon})$, where ω is the matrix multiplicative factor 2.37 and ϵ is the accuracy. Here the number of variables is d , so the running time is $O(d^{3.67} \cdot \log d \cdot \frac{1}{\epsilon} \cdot \log \frac{1}{\epsilon})$. As a result, solving the QP l times to get l directions will result in $O(l \cdot d^{3.67} \cdot \log d \cdot \frac{1}{\epsilon} \cdot \log \frac{1}{\epsilon})$. Calculating the skew on the l directions will take $O(l \cdot n \cdot d)$ time. The first part will be the dominant one in terms of time complexity.

In our technical report [16], we analyze the error bounds of the discretization approaches and provide the following negative result: let f' be a direction with a cosine similarity of (at least) $\cos(\alpha)$ with the best direction f^* . An upper-bound on skew ratio $\frac{f'}{f^*}$ depends on the dataset topology. As a result, when discretizing the directions with ϵ -angle spacing, *the error bound on the skew value depends on the topological characteristics of the dataset*. Still, we guarantee that for each direction f in the top- l high-skew local optima (Refer to Problem Definition 1), we evaluate a direction f' such that $\cos(f', f) \geq \epsilon$.

6.2 Exploration and Exploitation (E&E)

The E&E heuristic views the problem as a reinforcement learning instance and adapts the exploration-exploitation techniques for finding a set of high-skew directions. In particular, our goal is to select the next direction to evaluate sequentially to balance the explore-exploit trade-off. That is, starting with a small set of randomly selected directions, at each iteration, we either explore the space of possible directions or exploit the best possible directions found so far.

Without loss of generality, we adapt the epsilon-greedy approach [42]. Define $\epsilon \leq 1$ as the exploration probability. We flip a coin with probability ϵ in each iteration and choose to explore or exploit based on this trial result. Let f' be the best-known solution with the maximum skew so far while considering a hyper-parameter τ' for the cosine similarity of the exploitation phase. During each exploration step, we draw random directions in the space of all possible directions. On the other hand, during the exploitation, we select the best-known direction f' and select a random direction from the space of possible direction in the τ -vicinity of f' , defined as $\{f^{(s)} | \cos(f', f^{(s)}) \leq \tau'\}$. This problem can be reduced to the problem of sampling uniformly from the surface of a unit d -dimensional spherical cone with the center f' and the angular radius of τ . We use the sampling approach proposed in [8].

6.3 Focused Exploration

A significant challenge for the E&E approach is that in high dimensional cases where $d \gg 2$, the search space becomes huge, making it unlikely to find near-optimal directions. Therefore, instead, we first identify the under-performing regions with the hope of reducing the search space significantly. To do so, we use the model h_θ , trained on \mathcal{D} , to find a subset $S \subset \mathcal{D}$, which we call the *error region*, for which h_θ under-performs. Formally, the expected loss of the model in S is significantly higher compared to $\mathcal{D} \setminus S$.

After finding the error region, we focus on the directions that separate the error region S from other points $\mathcal{D} \setminus S$. In other words, we would like to find f such that after projecting \mathcal{D} on it, the error

Algorithm 2 Focused Exploration

```

1: procedure SMARTEXPLORE( $\mathcal{D}, h_\theta, \mathcal{K}, \mathcal{K}'$ )  $\triangleright h_\theta$  is the trained model on  $\mathcal{D}$ .
2:    $\mathcal{D}_{sorted} \leftarrow sort(\mathcal{D}, L_\theta)$   $\triangleright$  Sorts  $\mathcal{D}$  based on the loss function  $L_\theta$  (descending).
3:    $S \leftarrow$  First  $\mathcal{K}$  portion of  $\mathcal{D}_{sorted}$   $\triangleright S$  is the error region.
4:    $D^S \leftarrow sample(S)$   $\triangleright$  Down sample the error region uniformly.
5:    $D^{1S} \leftarrow sample(\mathcal{D} \setminus S)$ 
6:    $Heap \leftarrow []$   $\triangleright$  Initialize a max heap.
7:   for all vectors  $v \in \{s - t | s \in D^S, t \in D^{1S}\}$  do
8:      $f \leftarrow v / ||v||$ 
9:      $skew = GetSkew(\mathcal{D}, f)$ 
10:     $Heap.push(\langle f, skew \rangle)$ 
11:    $Output \leftarrow []$ 
12:   while  $Output$  size is less than  $l$  do
13:      $f \leftarrow Heap.pop()$ 
14:      $T \leftarrow p-tail(f)$ 
15:      $T_{err} \leftarrow$  First  $\mathcal{K}'$  portion of  $sort(T, L_\theta)$   $\triangleright$  Keep the challenging points of the tail
16:     if  $L_{T_{err}}(\theta) - L_{\mathcal{D}}(\theta) > \tau$  then  $Output.append(f)$ 
17:   return  $Output$ 

```

region is well-separated from others. We start by down-sampling S and $\mathcal{D} \setminus S$ to explore the specified directions faster.

Given the set of directions pushed into a Max-Heap (similar to Ray-sweeping), we start popping from the head of the heap and checking the most challenging points in the tail of these directions. Let $p-tail(f)$ be the tail of a popped direction f from the heap. For a parameter \mathcal{K}' , we check the top- \mathcal{K}' portion of $p-tail(f)$ to evaluate the condition of problem 1.

Many implementation details are deferred to § 7 for ease of explanation. A simplified pseudo-code of this approach is given in Algorithm 2. Based on our experiments, *this approach works well in practice on large-scale high-dimensional settings*.

7 EXPERIMENTS

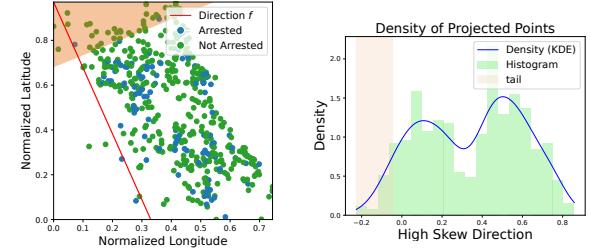
Having presented our algorithms in this section, we evaluate them experimentally for the 2D case and higher dimensions. We analyze our algorithms in each section and their effectiveness in finding the high-skew directions and candidate Minoria. The algorithms are implemented using Python, and the experiments are conducted on a workstation with Ubuntu 18.0 OS.

7.1 2D Experiments

We ran the 2D experiments on two real-world datasets: (a) **Chicago Crimes** dataset [1] and (b) **College admission** dataset [3].

7.1.1 Chicago Crimes. As stated in § 4, the 2D algorithms are applicable for datasets with more than two features as long as the projection directions are defined on a pair of attributes of interest. Therefore, in this experiment, we use the Chicago Crimes dataset [1] while using longitude and latitude for projection. This dataset contains the reported crimes in Chicago from 2001 until 2023. It has 22 columns and more than 1.4M rows.

Prediction Task. The target variable in this dataset is binary, reflecting whether a reported crime resulted in an arrest or not. We trained a simple Neural Network with one hidden layer of size 100 for this dataset. The model’s accuracy on the training dataset is 0.87, with an f1-score of 0.72.



(a) First high-skew direction. After projecting all the points.

Figure 4: Top high-skew directions on the Latitude and Longitude of Chicago Crimes dataset are presented in these two plots (the plots show a smaller sample of the dataset points). In the first plot, the output of the binary classification model is presented with green and blue lines. The red line shows the direction f with the highest projection skew. The orange region shows the tail region.

Hyper-parameters. In this problem, we first took a sample of size 1000 from the dataset and ran the Ray-sweeping algorithm on the two dimensions, Latitude and Longitude, to find the top high-skew directions with a minimum cosine similarity threshold $\cos(\frac{\pi}{12})$. We picked the top-3 high-skew directions. After selecting the top directions, we checked the tails of these directions in the original dataset (containing more than 1,400,000 rows). Running the Ray-sweeping algorithm in a sample of size 1000 takes around 15 seconds in the specified system configuration.

To search the entire space of directions, we first normalized the points to lie in the positive direction of the \mathbb{R}^2 space. In addition to running the Ray-sweeping from the x-axis to the y-axis, we also explored the negative scoped directions. To investigate the negative scoped directions, we transformed the points (x, y) in this space to $(y, A - x)$ for a large value A .⁷ This helps us run Ray-sweeping’s second pass on the transformed space to cover all negative directions.

Results. The top-2 high-skew directions are shown in Fig 4. As we can see, the first two high-skew directions are along the north side to the south side of Chicago. We select 0.01, 0.001, and 0.0001 percentile of the points in the tail of the high-skew projected points. The tails of the high-skew directions are presented using the orange regions. Evidently, in both cases, *the tails correspond to the North-side Chicago*. To provide some background about the city, Chicago is one of the most segregated cities in the US, with the vast majority of its north side being White.⁸ We refer the curious reader to [18, 30] for more details. North-side is under-represented in the dataset for many reasons, such as *predictive policing* [15]. At the same time, the f1-score significantly decreases on the tail of the projection, illustrating that predictions based on the crime dataset are not accurate for this group. To provide more details about these directions, we provide the tail-performance for different percentiles for these two directions in Tables 1 and 2. As illustrated in the tables, although the accuracy shows only a slight variation, there is a significant decrease in the F1-score on the tails in both cases.

⁷This can be viewed as rotating the points 90 degrees counter-clock-wise and renormalizing them.

⁸Please note that north-side being a Minoria (i.e., under-represented and under-performing) in the Crimes dataset is irrespective of its segregated nature.

Percentile	Accuracy	F1-score
1	0.87	0.72
0.1	0.87	0.62
0.01	0.89	0.68
0.001	0.85	0.40
0.0001	0.87	0.30

Table 1: (Chicago crime) Performance evaluation on the first high-skew direction (skew = 0.20). The tail is calculated on the entire dataset.

Percentile	Accuracy	F1-score
1	0.87	0.72
0.1	0.87	0.62
0.01	0.89	0.68
0.001	0.79	0.32
0.0001	0.81	0.18

Table 2: (Chicago crime) Performance evaluation on second high-skew direction (skew = 0.16). The tail is calculated on the entire dataset.

This indicates a pronounced class imbalance in the tail distribution, which is attributable to the demographic group imbalance.

7.1.2 College Admissions. The College-admission dataset [3] contains details about a group of students who sent applications to US universities. This dataset has 400 rows and 7 columns indicating the features of the submitted applications (like GRE, GPA, Gender/Race, etc.). The female-to-male ratio in this dataset is around 1.1.

Prediction Task. The prediction target is whether an application will result in an admission. We used a Logistic Regression model for this binary classification task. It gives us an accuracy of 0.70 and an f1-score of 0.36. We removed the sensitive attribute Gender from the dataset while training this classification model and never used Gender in any implementation steps.

Hyper-parameters. We run the 2D Ray-sweeping algorithm on the two columns GRE and GPA. We followed the same process explained for the Chicago crime dataset and investigated the tail of the highest-skew direction.

Results. In Fig 5a, we provide a visualization of GRE and GPA values from this dataset. The red line shows the highest skew direction. Fig 5b visualizes the tail of this high-skew direction (percentile $p = 0.1$). As we can see in Table 3, the ratio of the Female group is significantly larger in the 0.1, 0.08, and 0.04 percentiles compared to the whole dataset (1.1). The accuracy decreases up to 10 percent when concentrating on the tail of this direction. In this case, the Female group is not under-represented in the dataset because the ratio is almost 1.0. However, it is an under-performing and unknown minority group.

Runtime. The running time of the Ray-sweeping algorithm on the Chicago Crime dataset in our settings was 8 seconds. This time was 3.74 seconds for College-admission dataset.

7.1.3 Comparison with Baselines. We utilized two clustering techniques as baselines: k-means [28] and Affinity Propagation [23]. The resulting clusters were analyzed to identify potential candidates for the underrepresented demographic groups (Minoria). We employed k-means with various values of k and utilized Affinity Propagation with a carefully adjusted preference parameter to control the number of clusters and avoid excessive clustering. A

visualization of the k-means clusters on College-admission dataset is provided in Figure 5c.

The ratio of minority groups within each cluster in k-means for various datasets is presented in Table 4. As reflected in the table, unlike the region discovered using our methods (tail of skews), the ratios of minority groups within each cluster are not significantly different from the whole dataset (last row). As a result, these clusters are not good candidates for identifying *underrepresented* groups within the data. Furthermore, the number of points in each cluster is significantly larger compared to the regions we identified, making it more challenging to detect and analyze minority groups within such clusters. In contrast, the regions we discovered are more compact and therefore more efficient for mining *underrepresented* groups, offering a practical and focused approach to this task. Table 5 compares the results for various values of k , all yielding consistent results.

The results of applying Affinity Propagation clustering are presented in Figure 7b. Similar challenges persist, including an excessive number of clusters, requiring all clusters to be examined. Furthermore, the cluster ratios do not significantly differ from those observed in the entire dataset.

Runtime. The average run-time of k-means on the College admission dataset in our setting is 0.003 milliseconds. This value for the Affinity Propagation method is 0.13 seconds. Even though the running time is smaller for clustering methods compared to the Ray-sweeping algorithm, their output clusters are not valid for solving the Minoria Mining problem defined in this work.

7.2 Higher-dimensional Experiments

In this section, we evaluate the higher-dimensional methods proposed in § 6. As reflected in our negative result on the error-bound of the discretization approaches [16], their effectiveness depends on the dataset topology. Besides, their running time significantly increases by the number of dimensions. Our experiments verify these findings: we found discretization slow and less effective. Therefore, in the following we focus on the more practical approaches based on the search space exploration. Comparing the E&E method (§ 6.2) and the focused exploration approach (§ 6.3), our experiments found the latter (focused exploration) significantly more effective for Minoria mining. Therefore, in the following, we first present our results for this approach and then present our extended experiment results for the E&E approach and highlight some of its drawbacks.

7.2.1 Focused Exploration. We ran experiments on three real-world datasets. The **Adult** Income Dataset [20] consists of demographic information for individuals, from the 1994 Census database. The dataset is primarily used to predict whether a person earns more than \$50,000 per year based on various attributes (14 columns and 48,842 rows). Gender is the sensitive attribute in this dataset. The **COMPAS** dataset [7] contains data used to evaluate the effectiveness and fairness of the COMPAS recidivism prediction algorithm. This dataset includes criminal history, jail and prison time, demographics, and COMPAS risk scores for defendants from Broward County, Florida (52 columns and 18,316 rows). Race is the sensitive attribute in this dataset. The last dataset is the **Diabetes** dataset [40]. This is a well-known dataset used for binary classification tasks in the medical domain. It was collected by the National Institute of Diabetes and Digestive and Kidney Diseases and is available

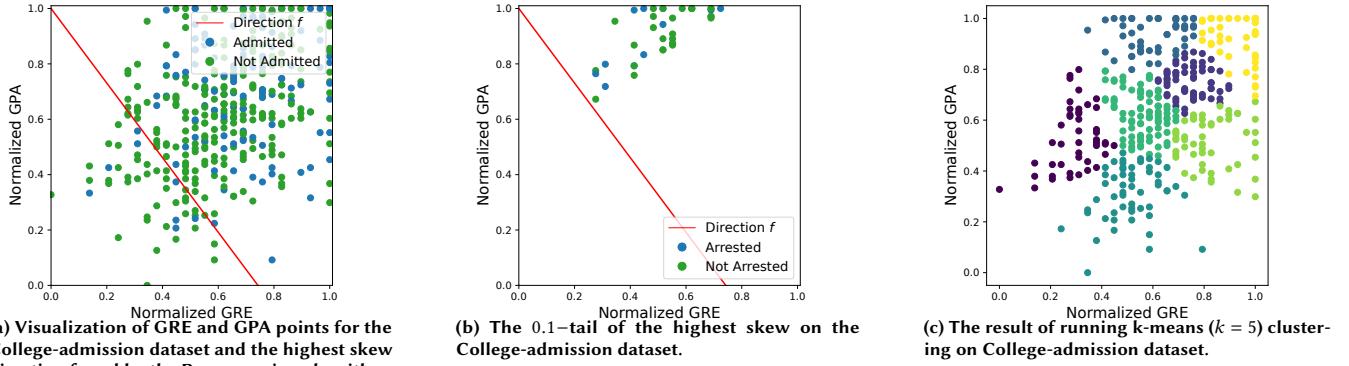


Figure 5: The result of running the Ray-sweeping algorithm on the College-admission dataset.

Percentile	Accuracy	F1 score	Female/Male ratio on Tail
1.00	0.70	0.36	1.10
0.50	0.68	0.42	1.12
0.20	0.67	0.48	0.80
0.10	0.61	0.34	2.00
0.08	0.64	0.42	1.81

Table 3: Evaluation on the first highest skewed direction (skew = 0.07) on College-admission dataset. The bolded values show a high ratio of minor groups in the tail.

through the UCI Machine Learning Repository (47 columns and 101,766 rows). Gender is the sensitive attribute in this dataset.

Prediction Task. For the Adult income dataset, we trained a Logistic Regression classifier to predict an individual’s income (whether it is more than \$50,000 or not) with an accuracy of 0.82 on the training data. For the COMPAS dataset, we trained a simple MLP (Multi-Layer Perceptron) classifier to predict the recidivism score of the individuals (Low, Medium, or High) with an accuracy of 0.71 on the training data. For the Diabetes dataset, we trained an MLP classifier to predict whether a patient is readmitted with an accuracy of 0.70 on the training data. We removed the sensitive attributes from these datasets before training the model, and we never used the information about sensitive attributes in any of our implementations. We only used the sensitive attributes to detect the effectiveness of final high-skew directions in finding Minoria.

Results. In these experiments, we hide the sensitive attribute containing the under-performing minority group and check if our mining experiments can discover it. To do so, we solved Problem 1 while setting $\ell = 5$. For the Adults and Diabetes datasets, we considered female as the minor demographic group. For the COMPAS dataset, we considered the race Caucasian (white) as the minor demographic group. The result of running the algorithm on the three datasets is shown in Fig 6. For the Adults dataset in Fig 6a, three of the five high-skew directions successfully show the Female group as the Minoria. This is because when focusing on the high error points in the tail of these directions, the ratio of Female points increases significantly. In addition, we can see that the ratio of the Female group in the tail of these high-skew directions is larger

Cluster ID	Size	Female/Male ratio
0	92	0.95
1	72	0.94
2	108	1.11
3	45	1.50
4	83	1.24
Total	400	1.10

Table 4: Evaluation of the output clusters of k-means clustering in detecting the unknown minority groups on College-admission dataset with $k = 5$.

k	Max Size	Maximum F/M ratio
2	201	1.26
3	136	1.20
4	104	1.20
5	108	1.50
6	90	1.44
7	93	1.50
Total	400	1.10

Table 5: Varying k in k-means on College-admission dataset. Col 2 shows the max Female to Male ratio on the output clusters.

than this ratio for the whole dataset. This shows that our model did not perform well on the minor group. For the COMPAS dataset in Fig 6b, two of the high-skew directions show the Race Caucasian as the Minoria. The red line is the high-skew direction that shows the model is significantly under-performing on this minor group. For the Diabetes dataset 6c, all the high-skew directions show that we have more error in the regions where the ratio of the minor group is larger than this ratio in the whole dataset. However, it is not easy for a human expert to detect this minor group as problematic because the ratio of this group is almost less than 0.3 in the high error subset of the tail. For the Diabetes dataset (Fig 6c), the tail skewed directions tend to show the problematic points as minor groups. As we focus on the top high error points of the tail of each direction, the ratio of the minor group increases. However, since this ratio is not large enough (less than 0.5), showing the points to a human supervisor may not determine the exact problematic feature in this region.

Runtime. The running time of Focused Exploration in our settings was 112 seconds on Adults dataset, 24 seconds on COMPAS, and 213 seconds on Diabetes dataset.

7.2.2 Comparison with Baseline. In this subsection, similar to subsection 7.1.3, we compare the results with two clustering algorithms: k-means and Affinity Propagation. Tables 6 and 8 show the ratio of minorities on the output clusters of k-means with $k = 5$. Same results as Section 7.1.3 is observed here. The result of comparing different values k is shown in Tables 7 and 9. Applying Affinity Propagation will also result in clusters which are not good candidates for Minoria mining (Figure 7a).

Runtime. The running time of k-means clustering on Adults and COMPAS datasets was 0.12 and 0.33 seconds, respectively. This running time for Affinity Propagation method was 14 and 20.2 seconds on Adults and COMPAS datasets. Even though the running time is smaller for clustering methods, their output clusters are not valid for solving the Minoria Mining problem defined in this work.

7.2.3 Extended Experiments: the E&E method. In this section, we present our experiment results for the E&E method discussed in § 6.2 on a synthetic dataset. We aim to evaluate the E&E method in *finding the high-skew directions* of a high-dimensional dataset in different scenarios on the demographic groups. We also mention some observed drawbacks of this method. We synthesized a dataset to explore the effectiveness of the E&E approach in finding high-skew directions. This dataset combines two normal distributions on \mathbb{R}^d (here, we used $d = 20$) with a total of 20,000 points. Each normal distribution plays the role of a demographic group (for example, male and female groups). The mean and the standard deviation of these two normal distributions are selected in a way that we can explore different scenarios in terms of how well the demographic groups can be separated linearly. After that, we down-sampled one of the regular groups to play the role of an under-represented group in the data. We removed any demographic information from this dataset and ran the E&E algorithm to find the high-skew directions.

Implementation. We started from 6 different random directions with an exploration probability of 0.4 ($\epsilon = 0.4$); otherwise, we exploited the best-found directions. Refer to § 6.2 for more details. We ran 10,000 iterations of E&E on this synthesized dataset. The running time, on average, was almost equal to 1 minute in the specified system configuration in the previous experiments.

Results. First, we explored a case when the demographic groups are well-separated linearly. This means there can be a linear classifier to classify these groups with an accuracy of more than %80. In this case, the E&E algorithm successfully finds the high-skew directions that separate the minor group from the major one. In Fig 8, the top-3 high-skew directions are found by this method. The plots show the histogram of the projected points on the high-skew direction f . In this case, by picking each direction's tail, we successfully target the minor group (red points). However, consider a case when the two demographic groups are hard to separate linearly (Fig 9). The top-3 high-skew directions do not help find the demographic groups in this case. The red points are almost overlapped with the major blue points. As a result, reporting the tail of this skew to a human expert does not help him to detect a Minoria.

The second challenge while using the E&E approach is the curse of dimensionality. It is hard to find the best highest-skew direction f^* in a high-dimensional case (where $d = 20$). As a result, it will take a long time for this method to converge when the size of a dataset is large (it took 1 minute for a dataset of size 20,000).

8 RELATED WORK

Addressing representation biases in datasets is a well-studied issue in the literature. Existing methods study the detection and resolution of representation bias for discrete and continuous data with various modalities [38]. The literature also includes statistical and classical methods for identifying patterns and gaps corresponding to empty spaces within the dataset [22, 26, 27].

Several studies have been conducted to identify the underperforming regions of machine learning models [14, 33, 35, 43]. These

studies aim to detect regions and sub-spaces that present challenges for the model in making accurate predictions. SliceFinder [14] detects specific “slices” of data where a machine learning model underperforms, indicated by a significantly higher average loss on these slices compared to other parts of the dataset. Other methods build on this concept by concentrating on specific regions of the data [35]. Chen et al. [12] employ clustering approaches to identify partitions of data that are discriminated against by a classification model. They aim to discover clusters where the model’s performance varies between demographic groups. The issue of identifying underrepresented groups within a sampled population is explored by Lancia et al. [25], who focus on detecting poorly represented groups and resampling these hard-to-sample sub-groups. Pastor et al. [33] also introduce a concept of pattern divergence to examine a model’s behavior across different sub-populations. Existing work also includes [31] that combines crowdsourcing and machine learning to detect uncovered groups over predetermined grouping attributes on image datasets without explicit group information.

The primary distinction between our work and the existing literature is that we assume no prior information about the grouping attributes within the dataset. In other words, we aim to mine Minoria from a collection of data points without explicit access to demographic group information over predetermined grouping attributes. Our goal is to identify the groups that are unknown, underrepresented, and underperforming for a given prediction task.

9 DISCUSSION

In this section, we discuss some aspects of our proposed approach and highlight its limitations.

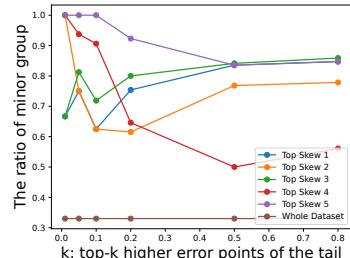
In general, our methods is limited to groups that are (almost) linearly separable. If the groups are linearly separable, the high-skewed direction found by our algorithms corresponds to a linear classification. In cases where linear separation is not feasible, one can consider introducing new features as polynomial combinations of existing features.

Not all the group candidates identified as the under-represented and under-performing regions in the data space are socially valid. Discovering such regions is intended to help model developers and data owners identify potential issues, decide if those are socially concerning, and *proactively* limit the scope of use of their datasets.

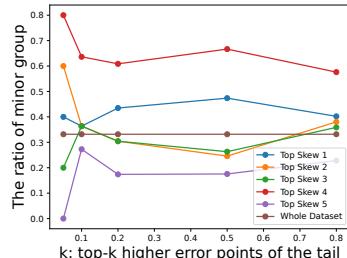
Our approach is only one of the possible solutions for Minoria mining. Finding alternative approaches requires extensive future research. Still, our experiments verified that clustering failed to identify the under-represented and under-represented groups.

Selecting appropriate columns from a dataset to perform Ray-sweeping is a critical practical decision. For instance, in our experiments, we utilized the *Latitude* and *Longitude* columns from the Chicago Crime dataset and the (GRE, GRA) pair from the College Admission dataset. One potential direction for column selection is identifying the most influential features for black-box machine learning models. Techniques such as feature importance measures or Shapley values can be employed to score each column, allowing the selection of the most impactful features for this algorithm.

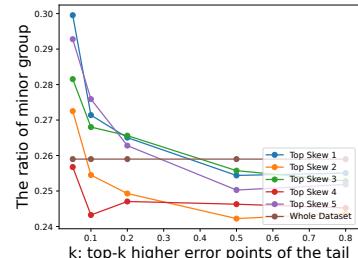
In higher dimensions, the curse of dimensionality introduces significant challenges. To address this, we developed heuristics aimed at exploring and exploiting these spaces more effectively.



(a) Adults dataset. Gender is the sensitive attribute.



(b) COMPAS dataset. Race is the sensitive attribute (Caucasian vs others)



(c) Diabetes dataset. Gender is the sensitive attribute.

Figure 6: Focused Exploration on the three datasets. Each line corresponds to one of the top-5 high-skew directions. The brown line shows the ratio of minor group in the whole dataset. The x-axis shows different values for a hyper-parameter p (or p -percentile of the tail). The ratio of minorities (y-axis) is calculated on the top- k (i.e. top p -percentile) high error points in the tail of each high-skew direction. If the ratio increases while decreasing p , we have more points from the minor group in the set of high error points of the tail of the specific direction.

Cluster ID	Size	female ratio
0	10850	0.13
1	5718	0.30
2	14678	0.52
3	3952	0.47
4	13644	0.24
Total	48842	0.33

Table 6: Evaluation of the output clusters of k-means when $k = 5$ algorithm in detecting the unknown minority groups for Adults dataset.

k	Max Size	Max female ratio
2	45252	0.33
3	19552	0.53
4	19522	0.52
5	14678	0.52
6	13041	0.53
7	13670	0.84
Total	48842	0.33

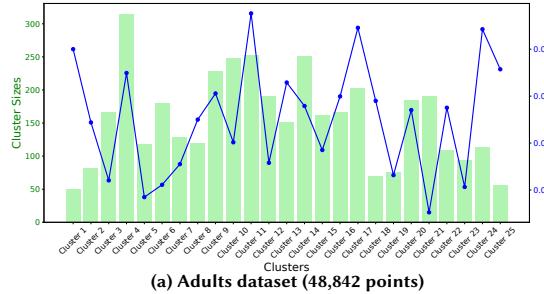
Table 7: Varying k in k-means for Adults dataset. Col 2 shows the max Female/total ratio in output clusters.

Cluster ID	Size	Caucasian ratio
0	524	0.09
1	5174	0.38
2	5860	0.29
3	2781	0.44
4	2680	0.25
Total	18316	0.33

Table 8: Evaluation of output clusters of k-means when $k = 5$ algorithm in detecting the unknown minority groups for COMPAS dataset.

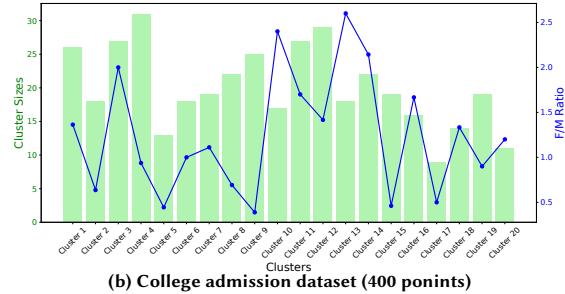
k	Max Size	Max Caucasian ratio
2	9020	0.35
3	7169	0.35
4	6638	0.45
5	5860	0.44
6	5100	0.46
7	4511	0.46
Total	18316	0.33

Table 9: Varying k in k-means for COMPAS dataset. Col 2 shows the max Caucasian/total ratio in output clusters.



(a) Adults dataset (48,842 points)

Figure 7: Affinity Propagation clustering: the size and the ratio of the minority group for each cluster are shown in different y-axes.



(b) College admission dataset (400 points)

Figure 8: Projected points on highest skewed directions discovered by E&E. In this case, the demographic groups (hidden during the learning phase) are well separated in the original dataset (the accuracy is almost 80%). The red points belong to the minor group.

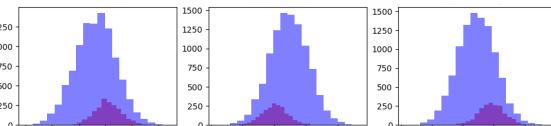


Figure 9: Projected points on highest skewed directions discovered by E&E. In this case, the demographic groups (hidden during the learning phase) are not linearly separable. The red points belong to the minor group.

These methods operate at a higher level of abstraction, focusing on identifying high-skew directions. Our findings demonstrate that these approaches, even though just heuristic, outperform traditional clustering methods in finding problematic regions within datasets.

10 CONCLUSION

In this paper, we introduced, formalized, and studied the problem of mining the unknown, underrepresented, and underperforming minority groups (Minoria). Making the connection to computational geometry concepts, we proposed efficient algorithms and practical solutions for finding high-skew projections as the combination of data features that reflect high-performance disparities at their tail. The scope of our proposed techniques is limited to the cases where the Minoria are linearly separable from the rest of the data. Extending our work to the more general class of non-linear projections, as well as developing alternative Minoria mining approaches, is an interesting direction for future work.

REFERENCES

[1] Chicago crimes - 2001 to present. <https://data.cityofchicago.org/Public-Safety/Crimes-2001-to-Present/ijzp-q8t2>. Accessed: 2024-05-20.

[2] City of chicago open data portal. <https://data.cityofchicago.org>.

[3] College admission data. <https://www.kaggle.com/datasets/eswarchandt/admission>. Accessed: 2024-05-20.

[4] Pankaj K Agarwal, Mark De Berg, Jiří Matoušek, and Otfried Schwarzkopf. Constructing levels in arrangements and higher order voronoi diagrams. In *Proceedings of the tenth annual symposium on Computational geometry*, pages 67–75, 1994.

[5] Noga Alon and Ervin Györi. The number of small semispaces of a finite set of points in the plane. *Journal of Combinatorial Theory, Series A*, 41(1):154–157, 1986.

[6] Nima Anari, Yang P Liu, and Thuy-Duong Vuong. Optimal sublinear sampling of spanning trees and determinantal point processes via average-case entropic independence. In *2022 IEEE 63rd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 123–134. IEEE, 2022.

[7] Julia Angwin, Jeff Larson, Surya Mattu, and Lauren Kirchner. Propublica compas dataset, 2016.

[8] Abolfazl Asudeh, HV Jagadish, Gerome Miklau, and Julia Stoyanovich. On obtaining stable rankings. *Proceedings of the VLDB Endowment*, 12(3):237–250, 2018.

[9] Abolfazl Asudeh, Azadeh Nazi, Nan Zhang, Gautam Das, and H.V. Jagadish. RRR: Rank-regret representative. In *SIGMOD*. ACM, 2019.

[10] Timothy M Chan. Output-sensitive results on convex hulls, extreme points, and related problems. In *Proceedings of the eleventh annual symposium on Computational geometry*, pages 10–19, 1995.

[11] Timothy M Chan. Remarks on k-level algorithms in the plane, 1999.

[12] Irene Chen, Fredrik D Johansson, and David Sontag. Why is my classifier discriminatory? *Advances in neural information processing systems*, 31, 2018.

[13] Sean Chester, Alex Thomo, S Venkatesh, and Sue Whitesides. Computing k-regret minimizing sets. *Proceedings of the VLDB Endowment*, 7(5):389–400, 2014.

[14] Yeounoh Chung, Tim Kraska, Neoklis Polyzotis, Ki Hyun Tae, and Steven Euijong Whang. Slice finder: Automated data slicing for model validation. In *2019 IEEE 35th International Conference on Data Engineering (ICDE)*, pages 1550–1553. IEEE, 2019.

[15] Andrea L DaViera, Marbella Uriostegui, Aaron Gottlieb, and Ogechi Onyeka. Risk, race, and predictive policing: A critical race theory analysis of the strategic subject list. *American journal of community psychology*, 73(1-2):91–103, 2024.

[16] Mohsen Dehghankar and Abolfazl Asudeh. Mining the minoria: Unknown, under-represented, and under-performing minority groups. *arXiv preprint arXiv:2411.04761*, 2024.

[17] Tamal K Dey. Improved bounds on planar k-sets and k-levels. In *Proceedings 38th Annual Symposium on Foundations of Computer Science*, pages 156–161. IEEE, 1997.

[18] James Diebel, Jacob Norda, and Orna Kretschmer. The demographic statistical atlas of the united states - statistical atlas. statisticalatlas.com/place/Illinois/Chicago/Race-and-Ethnicity#data-map/tract, [visited on June 2019].

[19] David P Doane and Lori E Seward. Measuring skewness: a forgotten statistic? *Journal of statistics education*, 19(2), 2011.

[20] Dheeru Dua and Casey Graff. Uci machine learning repository: Adult data set, 2019. Accessed: 2024-05-20.

[21] Herbert Edelsbrunner. *Algorithms in combinatorial geometry*, volume 10. Springer Science & Business Media, 1987.

[22] Jeff Edmonds, Jarek Gryz, Dongming Liang, and Renée J Miller. Mining for empty spaces in large data sets. *Theoretical Computer Science*, 296(3):435–452, 2003.

[23] Brendan J Frey and Delbert Dueck. Clustering by passing messages between data points. *science*, 315(5814):972–976, 2007.

[24] Francesco Gullo, Carlotta Domeniconi, and Andrea Tagarelli. Projective clustering ensembles. *Data Mining and Knowledge Discovery*, 26:452–511, 2013.

[25] C Lancia and E Riccomagno. Strategies for rare population detection and sampling: A methodological approach in liguria. *arXiv preprint arXiv:2405.01342*, 2024.

[26] Joseph Lemley, Filip Jagodzinski, and Razvan Andonie. Big holes in big data: A monte carlo algorithm for detecting large hyper-rectangles in high dimensional data. In *2016 IEEE 40th annual computer software and applications conference (COMPSAC)*, volume 1, pages 563–571. IEEE, 2016.

[27] Bing Liu, Liang-Ping Ku, and Wynne Hsu. Discovering interesting holes in data. In *IJCAI (2)*, pages 930–935, 1997.

[28] J Macqueen. Some methods for classification and analysis of multivariate observations. In *Proceedings of 5-th Berkeley Symposium on Mathematical Statistics and Probability/University of California Press*, 1967.

[29] Gabriela Moise, Arthur Zimek, Peer Kröger, Hans-Peter Kriegel, and Jörg Sander. Subspace and projected clustering: experimental evaluation and analysis. *Knowledge and Information Systems*, 21:299–326, 2009.

[30] Whet Moser. How redlining segregated chicago, and america. Chicago Magazine, Plotisc and City Life, 2017.

[31] Melika Mousaví, Nima Shahbazi, and Abolfazl Asudeh. Data coverage for detecting representation bias in image data sets: A crowdsourcing approach. In *Proceedings of the International Conference on Extending Database Technology (EDBT)*, 2024.

[32] Lance Parsons, Ehtesham Haque, and Huan Liu. Subspace clustering for high dimensional data: a review. *Acm sigkdd explorations newsletter*, 6(1):90–105, 2004.

[33] Eliana Pastor, Luca De Alfaro, and Elena Baralis. Looking for trouble: Analyzing classifier behavior via pattern divergence. In *Proceedings of the 2021 International Conference on Management of Data*, pages 1400–1412, 2021.

[34] Alex Rodriguez and Alessandro Laio. Clustering by fast search and find of density peaks. *science*, 344(6191):1492–1496, 2014.

[35] Svetlana Sagadeeva and Matthias Boehm. Sliceline: Fast, linear-algebra-based slice finding for ml model debugging. In *Proceedings of the 2021 International Conference on Management of Data*, pages 2290–2299, 2021.

[36] Nima Shahbazi and Abolfazl Asudeh. Reliability evaluation of individual predictions: a data-centric approach. *The VLDB Journal*, pages 1–28, 2024.

[37] Nima Shahbazi, Mahdi Erfanian, and Abolfazl Asudeh. Coverage-based data-centric approaches for responsible and trustworthy ai. *IEEE Data Eng. Bull.*, 47(1):3–17, 2024.

[38] Nima Shahbazi, Yin Lin, Abolfazl Asudeh, and HV Jagadish. Representation bias in data: a survey on identification and resolution techniques. *ACM Computing Surveys*, 55(13s):1–39, 2023.

[39] Micha Sharir, Shakhar Smorodinsky, and Gábor Tardos. An improved bound for k-sets in three dimensions. In *Proceedings of the sixteenth annual symposium on Computational geometry*, pages 43–49, 2000.

[40] JW. Smith, J.E. Everhart, W.C. Dickson, W.C. Knowler, and R.S. Johannes. Uci machine learning repository: Pima indians diabetes dataset, 1988.

[41] Chenkai Sun, Abolfazl Asudeh, HV Jagadish, Bill Howe, and Julia Stoyanovich. Mithralabel: Flexible dataset nutritional labels for responsible data science. In *Proceedings of the 28th ACM International Conference on Information and Knowledge Management (CIKM)*, 2019.

[42] Richard S Sutton and Andrew G Barto. *Reinforcement learning: An introduction*. MIT press, 2018.

[43] Ki Hyun Tae and Steven Euijong Whang. Slice tuner: A selective data acquisition framework for accurate and fair machine learning models. In *Proceedings of the 2021 International Conference on Management of Data*, pages 1771–1783, 2021.