

# Auditing for Spatial Fairness

Dimitris Sacharidis  
 dimitris.sacharidis@ulb.be  
 Université Libre de Bruxelles  
 Brussels, Belgium

George Papastefanatos  
 gpapas@athenarc.gr  
 IMSI/Athena Research Center  
 Athens, Greece

Giorgos Giannopoulos  
 giann@athenarc.gr  
 IMSI/Athena Research Center  
 Athens, Greece

Kostas Stefanidis  
 konstantinos.stefanidis@tuni.fi  
 Tampere University  
 Tampere, Finland

## ABSTRACT

This paper studies algorithmic fairness when the protected attribute is location. To handle protected attributes that are continuous, such as age or income, the standard approach is to discretize the domain into predefined groups, and compare algorithmic outcomes across groups. However, applying this idea to location raises concerns of gerrymandering and may introduce statistical bias. Prior work addresses these concerns but only for regularly spaced locations, while raising other issues, most notably its inability to discern regions that are likely to exhibit spatial unfairness. Similar to established notions of algorithmic fairness, we define spatial fairness as the statistical independence of outcomes from location. This translates into requiring that for each region of space, the distribution of outcomes is identical inside and outside the region. To allow for localized discrepancies in the distribution of outcomes, we compare how well two competing hypotheses explain the observed outcomes. The null hypothesis assumes spatial fairness, while the alternate allows different distributions inside and outside regions. Their goodness of fit is then assessed by a likelihood ratio test. If there is no significant difference in how well the two hypotheses explain the observed outcomes, we conclude that the algorithm is spatially fair.

## 1 INTRODUCTION

Algorithmic fairness refers to the notion that the algorithm (e.g., the machine, an ML model, an AI system) should not discriminate against individuals. Typically, discrimination is defined over groups of people that are considered *protected*, such as race or gender minorities. Abstractly, fairness requires that each population group, determined by a specific value to the *protected attribute* (e.g., race, gender), is on average treated or affected by the algorithm in the same manner. To make this requirement concrete, one first needs to define a *measure* to appropriately quantify the behavior, performance, etc. of the algorithm. Then, algorithmic fairness is achieved when the measure is statistically *independent* of the protected attribute. In practice, this mandates that the measure is distributed equally among protected groups. Fairness notions differ in how they define the measure. For example, statistical parity considers the positive rate (how often the algorithm assigns the positive/desirable class) as the measure, whereas equal opportunity [5] considers the true positive rate (how often the algorithm correctly assigns the positive/desirable class).

**Motivation.** In many cases, it is important to ensure that the algorithm does not discriminate against individuals on the basis of their location (place of origin, home address, etc.). That is, we consider location as the protected attribute and we want the algorithm to exhibit *spatial fairness*. For example, consider an algorithm that decides whether mortgage loan applications are accepted or not.<sup>1</sup> It is desirable that its decisions do not discriminate on the basis of the home address of the applicant. This could be to avoid *redlining*, i.e., indirectly discriminating based on ethnicity/race due to strong correlations between the home address and certain ethnic/racial groups, or to avoid *gentrification*, e.g., when applications in a poor urban area are systematically rejected to attract wealthier people. As another example, consider crime forecasting, where an algorithm predicts how likely a crime is to occur in a particular area. It is desirable that the algorithm is spatially fair in terms of its accuracy. That is, we require the predicted crime rate to not differ greatly than the observed crime rate *in all areas*. This could be to avoid *under- and over-policing*, and the sense of injustice they are typically associated with.


In cases like these, there is the common need for a principled method to *audit* an algorithm for spatial fairness, i.e., examine its outcomes and answer the question “*is it fair?*”. More importantly, if this answer is negative, the method should credibly *testify*, i.e., provide as evidence a region that is most likely to suffer from discrimination, and thus answer the question “*where is it unfair?*”. **Challenges.** Unlike typical protected attributes, such as race and gender, location is a continuous attribute. The group-based definition of fairness does not apply, in a straightforward manner, to *continuous protected attributes*. Rather, the standard approach is to first *discretize* the continuous domain to create groups, e.g., age or income groups, and then compare the outcomes for each group. The same idea can be applied to location, by defining non-overlapping spatial partitions (e.g., city blocks, zipcodes, districts), and computing the measure in each. This leads to a partitioning-based definition, according to which an algorithm is spatially fair if the measures per partition are equal.<sup>2</sup>

However, this simple partitioning-based definition has two drawbacks. Location is highly susceptible to *gerrymandering* [7, 17], which is the act of purposefully defining a partitioning of the space (via setting the partition boundaries, extents, shapes) so that the partition measures appear non-discriminatory. Moreover, conclusions drawn from comparing spatial aggregates highly depend on the shape and scale of the partitions (the areal units),

© 2023 Copyright held by the owner/author(s). Published in Proceedings of the 26th International Conference on Extending Database Technology (EDBT), 28th March-31st March, 2023, ISBN 978-3-89318-092-9 on OpenProceedings.org. Distribution of this paper is permitted under the terms of the Creative Commons license CC-by-nc-nd 4.0.

<sup>1</sup>In practice, an algorithm would compute a “credit score”, based on which a human would take the decision. In that case, we would like these scores to be spatially fair.

<sup>2</sup>A note on terminology: a partitioning of the space consists of a set of non-overlapping regions, called partitions, that collectively cover the space.



(a) Spatially fair-by-design outcomes distributed in Florida; each outcome has 0.5 probability of being positive. (b) Spatially unfair-by-design outcomes uniformly distributed; left half has twice more positive outcomes than right half.

Figure 1: Two spatial distributions each with  $n = 10,000$  outcomes among which  $p = 5,000$  are positive. According to [17], the fair-by-design distribution is less fair than the unfair-by-design distribution (MEANVAR of 0.0522 vs. 0.0431).

a well-known source of statistical bias termed the *modifiable areal unit problem* (MAUP) [17].

A better partitioning-based measure of spatial fairness that tries to address these two drawbacks is proposed in [17]. Briefly, the idea is to superimpose a high-resolution grid over the space, and consider all possible rectangular, grid-aligned partitionings of the space. In each partitioning, the variance of the measure in the partitions is computed. Then, the mean variance across all partitionings, hereafter denoted as MEANVAR, is computed. Lower values of MEANVAR suggest lower variance across the partitions in all partitionings and hence more fairness.

Note that MEANVAR is designed to assess the spatial fairness of outcomes that are *regularly distributed* in space: in each cell of the superimposed grid (or equivalently in each partition of the partitioning with the highest resolution) there is roughly the same number of outcomes. As a consequence, computing MEANVAR in the general case of outcomes arbitrarily distributed in space leads to counter-intuitive results. Concretely, MEANVAR cannot reliably audit an algorithm for spatial fairness. To illustrate this, consider the example presented in Figure 1. The algorithm producing the SemiSynth dataset depicted on the left is *spatially fair by design* as positive/desirable (resp. negative/undesirable) outcomes are randomly assigned to each location with a probability of 0.5, and are indicated as green (resp. red) points. The algorithm producing the Synth dataset depicted on the right is *spatially unfair by design*, as the left half of the area contains twice as many positive outcomes as the right half does. However, if we assess spatial fairness according to [17], we find that the unfair-by-design algorithm has a lower MEANVAR of 0.0431 and is thus considered more fair than the fair-by-design algorithm with a higher MEANVAR of 0.0522. This occurs because the spatial distribution of outcomes in Figure 1(a) is non-regular, and thus the number of outcomes within partitions varies greatly. Despite the fact that the ratio of positives is on average the same in each partition (by design), there exist several sparse partitions with few outcomes that are predominantly positive or negative, and which increase the variance of the measure. The example demonstrates that it is impossible to set a threshold for MEANVAR such that it distinguishes a fair from an unfair algorithm. Therefore, MEANVAR cannot answer the “*is it fair?*” question.



(a) A suspicious region in Iowa that makes the largest contribution to MEANVAR. (b) A region in California that testifies for spatial unfairness according to our framework (p-value < 0.005).

Figure 2: Regions most likely to be unfair according to different spatial fairness definitions for the LAR dataset described in Section 4.1. Depicted are the the number of observations ( $n$ ), the number of positive outcomes ( $p$ ), and the local positive rate ( $\rho$ ); the overall positive rate is 0.62.

Even if MEANVAR cannot discern fairness, one may wonder if it can be utilized to identify *suspicious regions*, i.e., regions that are likely to exhibit discrimination. The obvious way to search for suspicious regions is to consider the partitions that contribute the most to the MEANVAR value. These are the partitions whose measures are the farthest away from the partitioning mean and thus take extreme values. For the reasons discussed before, these partitions are likely to be sparse and predominantly positive or negative. An example is given in Figure 2(a), where the depicted partition with just five negative outcomes ties for the largest contribution to MEANVAR. On a first look, this appears to suggest an area of discrimination. However, this result is not statistically significant: it is not that uncommon to find a region that contains at least five negatives and no positives by chance (refer to [13]). Therefore, even though MEANVAR can identify regions with extreme measures, arguably it cannot meaningfully answer the “*where is it unfair?*” question.

**Our Solution.** In this work, we propose a practical definition for spatial fairness that enables *auditing* and *testifying*, i.e., that can answer both important questions regarding the algorithm, “*is it fair?*” and if not “*where is it unfair?*”. The definition applies to the most general case and does not assume regularly distributed observations.

Recall that algorithmic fairness is when some measure quantifying the behavior or performance of the algorithm is statistically independent of the protected attribute. In our setting, location is the protected attribute, so we naturally consider an algorithm to be *spatially fair*, if the measure is *independent of location*. This implies that for any region of the space, the distribution of the measure *inside* and *outside* the region should be the same.

To operationalize this definition, there are several challenges. The most complicating is how to determine the distribution of the measure within a region. If the region covers many observations, the observed (empirical) distribution of the measure is a good proxy for the actual distribution. Otherwise, what we observe in a sparse or small region might differ dramatically from what we observe outside it. Note that this issue does not manifest itself in categorical protected attributes (e.g., gender), simply because the number of observations per protected group (e.g., number of women and men) are typically very large.

To address this issue, instead of looking at the observed distribution, we want to express *how likely* it is to observe such a distribution if the algorithm was fair. Intuitively, we should expect to find a region with only four negative points clustered

together even when the algorithm is fair. Conversely, we should not expect to find a region that contains thirty negative points alone — observing such a region should be a stronger indication that the algorithm is not fair.

Inspired by the work in spatial scan statistics [6, 9, 16], we form two hypotheses and seek to quantify which one is a better fit for the data observed. The *null hypothesis* is that of spatial fairness, i.e., there is a single distribution that controls how the measure is distributed in the space, or mnemonically *inside = outside*. The *alternate hypothesis* states that there is a difference in the distribution inside and outside a region, or *inside  $\neq$  outside*. Given the observed data, we can determine the maximum likelihood for each hypothesis, compute their ratio, and test whether this likelihood ratio is statistically significant at a desired level.

The remainder of this paper is organized as follows. Section 2 reviews important concepts from algorithmic fairness and scan statistics. Section 3 introduces our definition of fairness and the auditing framework. Section 4 presents some results of our framework on several datasets. Section 5 concludes this work.

## 2 RELATED WORK

### 2.1 Algorithmic Fairness

Algorithmic fairness has seen various definitions [12]. One important distinction [2] is between *individual fairness* definitions are based on the premise that similar entities should be treated similarly, and *group fairness* definitions group entities based on the value of one or more protected attributes and ask that all groups are treated similarly. In this work, we consider the definition of the latter category. Note that fairness by unawareness, where the protected attribute is not considered by the algorithm is not sufficient in many cases, due to the presence of other attributes that might be correlated with them.

Let  $X$  denote a set of features that describe an individual, and let  $A$  denote the protected attribute(s); for ease of presentation assume that  $A$  is a binary valued attribute, where  $A = 1$  indicates the protected group. For what follows we assume that the algorithm is a binary classification model. Therefore, let  $Y$  denote the actual class of the individual (the ground truth), with the positive class  $Y = 1$  denoting the desirable/favorable class. For example, the positive class might correspond to the acceptance of a loan application, the hiring of a candidate, the granting of parole, etc. Moreover, let  $\hat{Y}$  denote the predicted output of the algorithm; note that although we do not explicitly show it,  $\hat{Y}$  depends on  $X$ .

Statistical approaches to group fairness can be distinguished as [4]: *base rates approaches* that use only the output  $\hat{Y}$  of the algorithm, and *accuracy approaches* that use both the output  $\hat{Y}$  of the algorithm and the ground truth  $Y$ . Base rate fairness compares the probability  $P(\hat{Y} = 1|A = 1)$  that one individual receives the favorable outcome when they belong to the protected group with the corresponding probability  $P(\hat{Y} = 1|A = 0)$  that one receives the favorable outcome when they belong to the non-protected group. To compare the two, we may take their ratio [3] or their difference [1]. When the probabilities of a favorable outcome are equal for the two groups, we have a special type of fairness termed *statistical parity*. Statistical parity preserves the input ratio, that is, the demographics of the individuals receiving a favorable outcome are the same as the demographics of the underlying population. Statistical parity is a natural way to model equity: members of each group have the same chance of receiving the favorable output.

Accuracy-based fairness warrants that various types of classification errors (e.g., true positives, false positives) are equal across groups. Depending on the type of classification errors considered, the achieved type of fairness takes different names [5]. For example, the case in which we ask that  $P(\hat{Y} = 1|Y = 1, A = 1) = P(\hat{Y} = 1|Y = 1, A = 0)$  (i.e., the case of equal true positive rate for the two groups) is called *equal opportunity*. Similarly, the case in which both the true positive rate and the false positive rate are equal for the two groups, is called *equal odds*.

These notions can be generalized under a common idea. Let  $M$  indicate the event/resource that is to be distributed fairly among individuals. In the case of statistical parity, the resource is positive rate and the event is  $M = \hat{Y}$ . In equal opportunity, the resource is true positive rate and thus the event is  $M = \hat{Y}|Y = 1$ . Equal odds considers the true positive rate and the false positive rate, which is modeled as  $M = \hat{Y}|Y = 0$ . Then these fairness notion can be expressed as the requirement that  $P(M|A = 1) = P(M|A = 0)$ , or equivalently that  $M$  is independent of  $A$ , i.e.,  $P(M|A) = P(M)$ .

More recently, there have been approaches that define fairness from a causal perspective [8, 10, 11]. For example, we might require that  $A$  does not *cause*  $M$ , rather than simply requiring that  $M$  is independent of  $A$ .

### 2.2 Spatial Fairness

The problem of spatial fairness has received little attention. As location is a continuous attribute, the standard approach is to discretize locations and apply the standard definitions [15]. This discretization approach may be susceptible to gerrymandering. To avoid this, [17] considers all possible grid-based rectangular partitionings of the space into regions. Given a partitioning, the algorithm is considered fair if it exhibits roughly the same performance in each region. This is quantified as the variance of some performance metric  $M$  across partitions. The algorithm is (perfectly) fair with respect to a partitioning if there is zero variance. The algorithm is (perfectly) fair if it is fair with respect to every partitioning, among a predefined set. In essence, this definition takes a brute-force approach to the problem: it applies the group-based approach, which is susceptible to gerrymandering, to each possible partition to eliminate the likelihood of gerrymandering.

[17] defines an unfairness measure as the mean variance across partitionings. This measure is then used as an additional optimization goal in a learning process to produce classification models that do not exhibit spatial unfairness in this sense. As discussed in the previous section, and as we illustrate in Section 4, the definition in [17] has several shortcomings making it less meaningful in real-life applications where the observations are not regularly distributed in space.

### 2.3 Scan Statistics

Our notion of spatial fairness is inspired by the work in *spatial scan statistics*. For a 1d domain like time, scan statistics answer questions such as whether there exists an unusually high concentration of events in some time period. In the spatial case, scan statistics answer similar questions, e.g., is there an unusually high concentration of a particular virus variant in some spatial area. Spatial scan statistics have been proposed for various underlying spatial processes, such as Bernoulli and Poisson [9], and multinomial distributions [6]. Our work is directly related to the former type. Moreover, there is another line of work on detecting *mixture areas* [14, 16], e.g., areas that contain a low or high diversity of point categories.

### 3 A SPATIAL FAIRNESS FRAMEWORK

We consider binary classification tasks concerning individuals. Let  $X$  denote a set of features that describe an individual. Among them, location  $L$  is considered the protected attribute. Let  $R$  denote a spatial region, and let  $n(R)$  denote the number of individuals whose location falls in the region. Moreover, let  $Y$  denote the true class of the individual, with the positive class  $Y = 1$  denoting the desirable class. Further, consider a classification model  $\hat{Y}$  that takes as input the features  $X$  of an individual and predicts their class  $\hat{Y}(X)$ . A summary of the important notation is in Table 1.

In what follows, we define spatial fairness in terms of statistical parity, i.e., we quantify the positive rate of the model; the definitions can be adapted to the cases of equal opportunity and equal odds by replacing  $\hat{Y}$  with  $\hat{Y}|Y = 1$  to account for the true positive and  $\hat{Y}|Y = 0$  for the false positive rates.

Let  $\rho = Pr(\hat{Y} = 1)$  denote the *positive rate* (*pr*) of the model, i.e., the probability that the models predicts the positive class. Restricting focus to a particular spatial region  $R$ , we define the *local positive rate* as  $\rho(R) = Pr(\hat{Y} = 1|L \in R)$ , which is the *pr* taking into account only individuals within the region  $R$ .

We now formulate an *idealized* notion of spatial statistical parity, by requiring that every local positive rate is equal to the positive rate of the model, i.e.,  $\rho(R) = \rho$  for all  $R$ ; equivalently, all local positive rates should be equal. This definition is impractical as it can only be satisfied by trivial classifiers that predict either 0 or 1. Thus, we must relax the definition to allow the local positive rates to differ from each other and the overall *pr*. The question is to what extent they can differ before we declare the model to be spatially unfair. To answer this, we define a statistical test, based on the Bernoulli spatial scan statistic [9], that decides whether the positive rate is *homogeneous* across the space.

First, observe that we can interpret a positive rate  $\rho$  as the probability that the model assigns an individual to the positive class; i.e., a Bernoulli trial with success probability  $\rho$ . Now, consider the group of people within a region  $R$ . The number of positive labels the model assigns follows the Binomial distribution with  $n(R)$  trials and success probability  $\rho$ , denoted as  $p(R) \sim B(n(R), \rho)$ .

We now define two hypotheses and use a statistical test to determine which explains the observed data better. The discussion that follows is based on the multinomial spatial scan statistic; an important difference is that we do not care for the direction of change of the statistic inside and outside a region. The null hypothesis  $\mathcal{H}_0$  states that in every region  $R$  the number of positive labels,  $p(R)$ , follows the Binomial distribution  $B(n(R), \rho_0)$ , where  $R, \rho_0$  are the parameters of the null hypothesis. The alternative hypothesis  $\mathcal{H}_1$  states that there is a region  $R$  such that  $p(R)$  follows  $B(n(R), \rho_0)$ , while the number of positive labels outside  $R$  follows a Binomial with a different success probability  $\rho_1 \neq \rho_0$ , i.e.,  $B(N - n(R), \rho_1)$ . The alternate hypothesis has three parameters  $R, \rho_0, \rho_1$ .

To quantify which hypothesis explains the observed data better, we derive their maximum likelihoods, and compute the likelihood ratio. Let us first consider the null hypothesis. The likelihood of  $\mathcal{H}_0$  is given by  $L_0(R, \rho_0) = \rho_0^{p(R)} (1 - \rho_0)^{n(R) - p(R)}$ , where  $n(R)$  (resp.  $p(R)$ ) denotes the number of individuals (resp. with positive labels) in a region  $R$ . For any region  $R$ , observe that the maximum likelihood is when  $\rho_0 = p(R)/n(R)$ . Across all regions, the likelihood takes the maximum value when  $R$  is the entire space. Therefore, the maximum likelihood value of the null hypothesis is  $L_0^{max} = \rho^P (1 - \rho)^{N - P}$ .

**Table 1: Important Notation**

Symbol	Meaning
$X$	features of an individual
$L$	location of an individual
$Y$	actual outcome/class
$\hat{Y}$	model prediction
$N$	number of individuals
$P$	num. of indiv. predicted to be in the positive class
$\rho = \frac{N}{P}$	model's positive rate
$R$	a spatial region
$n(R)$	number of individuals in $R$
$p(R)$	num. of indiv. in $R$ predicted to be in the pos. class
$\rho(R) = \frac{n(R)}{p(R)}$	model's local positive rate in $R$

Consider now the alternative hypothesis. Its likelihood is the product of two binomials, one for inside and one for outside the region  $R$ , each having a different success probability:

$$L_1(R, \rho_0, \rho_1) = \rho_0^{p(R)} (1 - \rho_0)^{n(R) - p(R)} \rho_1^{P - p(R)} (1 - \rho_1)^{N - n(R) - (P - p(R))}$$

First, we maximize the likelihood for a given  $R$ . Recall that  $\mathcal{H}_1$  requires that  $\rho_0 \neq \rho_1$ . If the observed positive rate inside the region,  $p(R)/n(R)$ , is different from that outside the region,  $(P - p(R))/(N - n(R))$ , then the likelihood takes its maximum value when  $\rho_0$  and  $\rho_1$  take the values of the observed positive rates inside and outside the region, respectively. Otherwise, it cannot exceed the likelihood when  $\rho_0 = \rho_1 = \rho$ . Concretely, we have:

$$L_1^{max}(R) = \begin{cases} L_1(R, \frac{p(R)}{n(R)}, \frac{P - p(R)}{N - n(R)}) & \text{if } \frac{p(R)}{n(R)} \neq \frac{P - p(R)}{N - n(R)}, \\ L_0^{max} & \text{otherwise.} \end{cases} \quad (1)$$

We refer to the value in Eq. 1 as the *spatial unfairness likelihood*, and denote it as SUL. The next step is to identify the region that maximizes Eq. 1. This is achieved by going over a predetermined set of regions  $\mathcal{R}$ . Let  $R^* \in \mathcal{R}$  be the region that maximizes Eq. 1. We compute the likelihood ratio test statistic of the alternative over the null hypothesis as  $\tau = \frac{L_1^{max}(R^*)}{L_0^{max}}$ . Note, that in practice, we compute and determine the difference of log-likelihoods.

To determine how significant the test statistic  $\tau$  value is, we need to perform a Monte Carlo simulation to determine its distribution, as also performed in [9, 16]. Specifically, the simulation goes as follows. We create alternate worlds assuming that the  $N$  individuals are located as in our data, but their label is determined by a Bernoulli trial with success probability  $\rho$ . For each alternate world, we compute the  $\tau$  statistic. Suppose we simulate  $w - 1$  worlds, and the  $\tau$  statistic of the real world ranks at the  $k$ -th highest position among all worlds. Then, the  $p$ -value of the real world's statistic is  $k/w$ .

We declare an algorithm as *spatially fair* if the  $p$ -value of the aforementioned likelihood ratio test statistic is below some predefined significance level  $\alpha$ . In this case, the observed data are just as likely to be generated under the spatial fairness null hypothesis as under the alternate hypothesis. Otherwise, we have to reject the spatial fairness hypothesis. In this case, we may offer evidence that support the hypothesis that the algorithm is unfair. Providing evidence is the *identification* process. We consider all examined regions that have a statistically significant likelihood ratio, and we rank them in decreasing order of their likelihood ratio. We then return the top- $k$  regions as evidence.

The computational complexity of our auditing framework is  $O(M \cdot N \cdot Q)$ , where  $N$  is the number of regions to scan,  $Q$  the average cost of a spatial range-count query, and  $M - 1$  is the number of Monte Carlo simulations.

## 4 EXPERIMENTS

Our experimental study investigates two research questions. R1. In a setting compatible with prior work [17], does our spatial fairness definition present more meaningful results compared to prior work? R2. Does our method identify regions that can be considered as potentially exhibiting spatial unfairness? Section 4.1 describes the dataset used, while Sections 4.2 and 4.3 answer the two research questions.

The code and datasets are available in GitHub.<sup>3</sup>

### 4.1 Datasets

In our experiments, we use two real, one semi-synthetic, and one synthetic datasets. The first real dataset concerns about mortgage loan applications. Thanks to the Home Mortgage Disclosure Act, financial institutions in the US are required to publicly disclose information about mortgages. We download the modified loan/application register (LAR) records<sup>4</sup> for Bank of America for the year 2021. The dataset contains information about over 300 thousand applications, including whether they were denied and the census tract of the applicant, which is a geographic region roughly corresponding to cities and towns. We preprocess the dataset to keep records for loan applications approved or denied. Further, we use the Gazetteer files provided by the US Census Bureau<sup>5</sup> to associate each census tract with the coordinates of its center. In the end, we obtain a dataset, denoted as LAR, with 206,418 applications, among which 127,286 were granted (i.e., 0.62 positive rate), geographically distributed along 50,647 locations (refer to [13] for a depiction). We use LAR to audit spatial fairness when the measure of interest is the positive rate, i.e., in the *statistical parity* sense. That is we want to investigate if all areas have the same chance of being granted a mortgage loan. The overall positive rate is 0.62.

The second real dataset is about crime incidents in the city of Los Angeles from 2010–2019<sup>6</sup>. The data contains several attributes including an approximate location of the incident. We consider the incident code, which we binarize into serious and non-serious crimes, as the label. We train a random forest classifier to predict the “seriousness” of the incident; the positive class is the serious crimes, while the negative the non-serious. As features, we select 7 attributes: time, code of police precinct, victim’s age, sex and descent, the type of the incident location, and the weapon used. After removing entries with missing values, there are 711,852 incidents split into train/test sets with a 70%:30% ratio. We collect the predictions and the true labels for the incidents in the test set, and construct our dataset, denoted as Crime. The accuracy of the model is 0.78. We use Crime to audit spatial fairness where the measure of interest is the true positive rate. That is, in the *equality of opportunity* sense, we want to investigate if the algorithm’s accuracy for serious crimes is independent of location. Therefore, to apply the framework of [17] and ours, we retain the predictions for the true positive labels, which results in 61,266 entries, and the overall true positive rate is 0.58 (refer to [13] for a depiction).

The semi-synthetic dataset, denoted as SemiSynth and depicted in Figure 1(a), contains 10,000 outcomes for locations that are randomly selected in Florida from the LAR dataset. The positive and negative are randomly assigned to each location with a

probability of 0.5. Hence, SemiSynth is *spatially fair by design*. The measure of interest is positive rate.

The synthetic dataset, denoted as Synth and depicted in Figure 1(b), contains 10,000 outcomes for locations selected uniformly at random within a rectangular area. The area is split into two halves, each containing 5,000 outcomes. However, the left half has twice as many positive outcomes as the right half does. Therefore, the positive rate in the left half is about 0.67, while in the right half is 0.33.

### 4.2 Results on Partitionings

In the first round of experiments, we consider a setting appropriate for the MEANVAR measure of spatial unfairness proposed in [17]. Therefore, we consider rectangular partitionings of the space. To compare our methodology with MEANVAR, we restrict our methodology to only audit for fairness the partitions that belong to the partitionings.

**Is it Fair?** The first experiment is to test whether the two methods can correctly audit for fairness, i.e., answer the “*is it fair*” question. For this we use the SemiSynth and Synth datasets that we have designed to be fair and unfair, respectively. We construct 100 rectangular partitionings, where the number of horizontal and vertical splits of the space is randomly selected between 10 to 40.

As discussed in the context of Figure 1, MEANVAR is 0.0522 for the fair-by-design SemiSynth and 0.0431 for the unfair-by-design Synth; recall that lower MEANVAR values suggest more fairness. Therefore, MEANVAR fails to discern fairness. In contrast, our method finds that SemiSynth is fair while Synth is unfair at the 0.005 statistical significance level.

**Where is it Unfair?** In the second set of experiments, we partition the space into a regular grid of various granularities, and investigate whether the methods can meaningfully identify partitions that potentially exhibit spatial unfairness. For MEANVAR, we identify the partitions that make the largest contribution to the MEANVAR value. As discussed, in Section 1, these are the partitions with extreme measures that differ greatly from the partitioning mean. We rank partitions according to their contribution to MEANVAR. For our method, we fix the statistical significance level at 0.005, and identify all partitions that exhibit spatial unfairness at that significance level. We rank partitions according to their SUL value (Eq. 1).

We first examine spatial fairness in terms of positive rate (i.e., statistical parity) and use the LAR dataset. Intuitively, we want to identify partitions where the algorithm assigns the positive/desirable class differently from the global mean. We consider a partitioning of high resolution  $100 \times 50$ . Our framework declares the outcomes as spatially unfair and identifies 59 statistically significant partitions depicted in Figure 3(a). The partition that has the highest SUL value is displayed in Figure 2(b), and concerns a region in northern California that covers almost 8,000 outcomes, among which 84% are positive. Our framework identifies mostly dense regions that have statistically significant different local positive rates compared to the global. In contrast, we display the top-50 partitions according to MEANVAR in Figure 3(b). It is easy to observe that they are all very sparse partitions that contain only negative outcomes. The largest of them with 5 outcomes is displayed in Figure 2(a).

Comparing the results of highest MEANVAR and SUL in Figure 2, we find that the partition in California has a non-extreme local positive of 0.84, compared to 0 for the partition in Iowa, whereas the former has a much higher and significant log-likelihood

<sup>3</sup><https://github.com/dsachar/AuditSpatialFairness>

<sup>4</sup><https://ffiec.cfbp.gov/data-publication/modified-lar/2021>

<sup>5</sup><https://www.census.gov/geographies/reference-files/time-series/geo/gazetteer-files.2021.html>

<sup>6</sup><https://data.lacity.org/Public-Safety/Crime-Data-from-2010-to-2019/63jg-8b9z>



(a) Spatial Fairness: the 59 statistically significant unfair partitions



(b) MEANVAR-Based Fairness: top-50 partitions of highest MEANVAR

Figure 3: LAR: Results for a high-resolution partitioning of  $100 \times 50$ .

difference of about 1000 compared to 0.96 for the latter (in this experiment log-likelihood differences beyond 9.6 are significant at the 0.005 level). As expected, MEANVAR identifies partitions with extreme measures, whereas SUL identifies partitions that have abnormal (in the statistical likelihood sense) measures.

We now examine spatial fairness in terms of true positive rate (i.e., equal opportunity) and use the Crime dataset. Intuitively, we want to identify partitions where the prediction accuracy is different from the overall prediction accuracy of the algorithm. We consider a partitioning of low resolution  $20 \times 20$ . Our framework declares the outcomes as spatially unfair and identifies 5 statistically significant partitions depicted in Figure 4(a). One of the partitions with the highest SUL value is located in Hollywood, and covers almost 3,000 outcomes, where only 51% of them are predicted to be serious crimes. Contrasting this to the global true positive rate of 0.58, we see that the algorithm tends to classify incidents as non-serious in this area, more so than in other areas. In contrast, we display the top-5 partitions according to MEANVAR in Figure 4(b). All of them concern very sparse areas with a single false positive, and are thus not interesting for the auditor.

### 4.3 Results for Unrestricted Regions

In this set of experiments, we only study our framework and can thus examine arbitrarily sized regions. Specifically, we consider square regions with 20 different side lengths ranging from 0.1 up to 2 degrees (roughly 10 to 200 kilometers). The centers of square regions are placed in 100 locations defined as the centers of a k-means clustering of the observation locations. In total, we scan 2,000 square regions. Their centers and their smallest and largest shape are depicted in [13].

Applying our framework at a significance level of 0.005, we identify 700 unfair regions. As these regions intersect each other, we select a set of non-overlapping regions. We examine centers in sequence, and for each center we keep the region with the highest value of the statistic. Fig 5 displays the 28 non-overlapping regions. The important observation is that our framework identifies regions of varying area size and observation size. For example an area near Tampa, FL is the smallest radius of 0.1 degrees with



(a) Spatial Fairness: the 5 statistically significant unfair partitions



(b) MEANVAR-Based Fairness: top-5 partitions of highest MEANVAR

Figure 4: Crime: Results for a partitioning of  $20 \times 20$ .



Figure 5: LAR: The 28 non-overlapping unfair regions

the largest number, 473, of observations, while a nearby area centered in Orlando, FL is the largest, with a radius of 1 degrees, and 4,783 observations.

In conclusion, the regions returned as evidence from our methodology that the algorithm is unfair are non-trivial (especially, compared to those discovered by MEANVAR), and can be used from the auditor to further investigate whether the observed unfairness is justified or not.

## 5 CONCLUSION

This paper introduces a generally applicable definition of spatial fairness. It includes a concrete framework to audit an algorithm for spatial fairness, and the ability to identify regions that are very likely to be spatially unfair with statistical significance. The intuition is that for any region, the distribution of outcomes inside and outside the region should be similarly distributed. The statistical test examines if the spatial fairness assumption is more likely than the alternative that allows for regions with different outcome distributions inside and outside.

**Acknowledgement:** The authors were partially supported by the EU’s Horizon Programme call, under Grant Agreements No. 101093164 (ExtremeXP) and No. 101070568 (AutoFair).

## REFERENCES

- [1] Toon Calders and Sicco Verwer. 2010. Three naive Bayes approaches for discrimination-free classification. *Data Min. Knowl. Discov.* 21, 2 (2010), 277–292.
- [2] Cynthia Dwork, Moritz Hardt, Toniann Pitassi, Omer Reingold, and Richard S. Zemel. 2012. Fairness through awareness. In *Innovations in Theoretical Computer Science 2012*. ACM, 214–226.
- [3] Michael Feldman, Sorelle A. Friedler, John Moeller, Carlos Scheidegger, and Suresh Venkatasubramanian. 2015. Certifying and Removing Disparate Impact. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. ACM, 259–268.
- [4] Sorelle A. Friedler, Carlos Scheidegger, Suresh Venkatasubramanian, Sonam Choudhary, Evan P. Hamilton, and Derek Roth. 2019. A comparative study of fairness-enhancing interventions in machine learning. In *Proceedings of the Conference on Fairness, Accountability, and Transparency, FAT\* 2019*. ACM, 329–338.
- [5] Moritz Hardt, Eric Price, and Nati Srebro. 2016. Equality of Opportunity in Supervised Learning. In *NIPS*. 3315–3323.
- [6] Inkyung Jung, Martin Kulldorff, and Otsukei John Richard. 2010. A spatial scan statistic for multinomial data. *Statistics in medicine* 29, 18 (2010), 1910–1918.
- [7] Michael Kearns, Seth Neel, Aaron Roth, and Zhiwei Steven Wu. 2018. Preventing Fairness Gerrymandering: Auditing and Learning for Subgroup Fairness. In *Proceedings of the 35th International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 80)*, Jennifer Dy and Andreas Krause (Eds.). PMLR, 2564–2572. <https://proceedings.mlr.press/v80/kearns18a.html>
- [8] Niki Kilbertus, Mateo Rojas-Carulla, Giambattista Parascandolo, Moritz Hardt, Dominik Janzing, and Bernhard Schölkopf. 2017. Avoiding Discrimination through Causal Reasoning. In *NIPS*. 656–666.
- [9] Martin Kulldorff. 1997. A spatial scan statistic. *Communications in Statistics-Theory and methods* 26, 6 (1997), 1481–1496.
- [10] Matt J. Kusner, Joshua R. Loftus, Chris Russell, and Ricardo Silva. 2017. Counterfactual Fairness. In *NIPS*. 4066–4076.
- [11] Razieh Nabi and Ilya Shpitser. 2018. Fair Inference on Outcomes. In *AAAI*. AAAI Press, 1931–1940.
- [12] Evaggelia Pitoura, Kostas Stefanidis, and Georgia Koutrika. 2022. Fairness in rankings and recommendations: an overview. *VLDB J.* 31, 3 (2022), 431–458.
- [13] Dimitris Sacharidis, Giorgos Giannopoulos, George Papastefanatos, and Kostas Stefanidis. 2023. Auditing for Spatial Fairness. <https://doi.org/10.48550/ARXIV.2302.12333>
- [14] Dimitrios Skoutas, Dimitris Sacharidis, and Kostas Patroumpas. 2021. Discovering Mixture-Based Best Regions of Arbitrary Shapes. In *Proceedings of the 29th International Conference on Advances in Geographic Information Systems*. 468–479.
- [15] Leonard Weydemann, Dimitris Sacharidis, and Hannes Werthner. 2019. Defining and measuring fairness in location recommendations. In *Local-Rec@SIGSPATIAL*. ACM, 6:1–6:8.
- [16] Yiqun Xie, Han Bao, Yan Li, and Shashi Shekhar. 2020. Discovering spatial mixture patterns of interest. In *Proceedings of the 28th International Conference on Advances in Geographic Information Systems*. 608–617.
- [17] Yiqun Xie, Erhu He, Xiaowei Jia, Weiye Chen, Sergii Skakun, Han Bao, Zhe Jiang, Rahul Ghosh, and Praveen Ravirathinam. 2022. Fairness by "Where": A Statistically-Robust and Model-Agnostic Bi-level Learning Framework. In *AAAI*. AAAI Press, 12208–12216.