2023 IEEE International Conference on Robotics and Automation (ICRA 2023)
May 29 - June 2, 2023. London, UK

Socially Fair Coverage Control

Matthew Malencia, George Pappas, Vijay Kumar

***Abstract*—We investigate and develop algorithms for social fairness in coverage control problems. Existing coverage control methods are efficient, optimizing the average expected distance from any event to the nearest robot. However, in societal applications like disaster response or transportation, these conventional objectives lead to disparate coverage costs with respect to different groups within a population. We formulate social fairness for coverage control as the minimization of the maximum coverage cost among a set of groups within a population. Our approach uses Voronoi iteration to solve this novel problem by approximating the non-differentiable objective with the log-sum-exp and defining a gradient based controller that prioritizes fairness while also optimizing average performance when disparities between groups are low. We show convergence properties of this proposed control law and demonstrate the approach in simulations of randomly generated population densities as well as environments generated from U.S. census data on population rates and demographics. Our approach provides greater fairness than existing methods while maintaining similar computational time and convergence properties.**

***Index Terms*—Fairness, Distributed Robot Systems, Multi Robot Systems, Robustness, Ethics and philosophy**

I. Introduction

Multi robot systems are becoming increasingly present in society. Future robotics systems will exist not only in our factories, but in our cities and our daily lives. Because these systems will be applied *within* social contexts, fairness is a crucial consideration to ensure that the robotics systems are both trusted and legal.

A common application of multi robot systems is distributed mobile sensing, or coverage control, where a team of robots are distributed over a region to detect events of interest. Coverage control is used in environmental monitoring and sensing [1], surveillance [2], and urban transportation [3]. The objective of coverage control is to position a robotic sensor network to maximize the likelihood of detecting events [4].

Many of the applications of coverage control are societal, like surveillance and transportation, where events are directly tied to people, e.g., those hailing a ride, or natural disaster victims seeking help. Therefore, fairness is a necessary characteristic for these systems to be trusted by those being served. To highlight how such systems can create disparate outcomes, consider the natural disaster scenario in Example 1, as depicted in Figure 1.

Example 1 (Surveillance)**.** *An emergency response team is deployed to Philadelphia County, Pennsylvania, USA, after a natural disaster. First responders can use a team of quadrotors to surveil the area for people in need of help. The robots use a population density map to collectively optimize the likelihood of detecting those in need of help. Naturally, the robots stay in densely populated areas. However, groups within a population (e.g., racial groups) are not uniformly distributed among the population. Therefore, the likelihood of missing someone can vary drastically between groups. Clearly,* ***this approach is unfair*** *with respect to the groups. Ideally,* ***the likelihood of detection should be independent of group status*** *such as socio-economic status, age, or sensitive characteristics like race, religion, and ethnicity.*

Contributions. We define for the first time social fairness with respect to sensitive characteristics for coverage control applications. Our proposed solution approximates the socially fair coverage objective and defines a gradient based controller for Voronoi iteration. Our main contributions are:

- We show the disparate impact that can occur with respect to groups within a population when deploying multi robot systems using traditional coverage control techniques.
- We define the Socially Fair Coverage Problem as the minimization of the maximum coverage cost among a set of groups within a population.
- We propose a gradient based controller for Voronoi iteration using a smooth approximation of the objective function and solve the Socially Fair Coverage Problem with convergence guarantees.

We provide empirical analysis of our approach using simulations of random environments and of real U.S. census data. We show that our approach provides fair coverage irrespective of disparate population rates between majority and minority groups. Lastly, our approach improves fairness on real population data with no increase in computation or convergence time compared to traditional approaches.

II. Background

A. Coverage Control Preliminaries

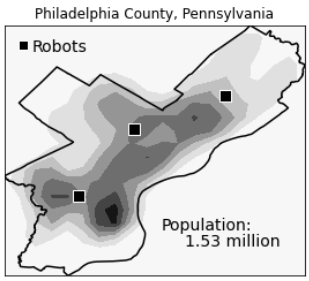
Coverage control seeks to distribute a team of robots over a compact and convex spatial region $\mathcal{D} \subset \mathbb{R}^2$ to detect events, e.g., people in need of help (Example 1), where the density function $\phi : \mathcal{D} \to \mathbb{R}^+$ represents event likelihood.

We gratefully acknowledge the support of ARL DCIST CRA W911NF-17-2-0181, NSF Grants CCR-2112665, and EEC-1941529, ONR grant N00014-20-S-B001, Qualcomm Research, NIFA grant 2022-67021-36856, the IoT4Ag Engineering Research Center funded by the National Science Foundation (NSF) under NSF Cooperative Agreement Number EEC-1941529. The first author acknowledges support from the National Science Foundation Graduate Research Fellowship under Grant No. DGE-1845298.

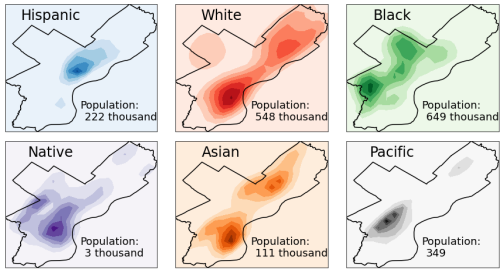
The authors are with GRASP Laboratory, Levine Hall 4th floor, University of Pennsylvania, 3330 Walnut Street, Philadelphia, PA 19104-6228. *corresponding author's email:* {malencia}@seas.upenn.edu

2023 IEEE International Conference on Robotics and Automation (ICRA) | 979-8-3503-2365-8/23/$31.00 ©2023 IEEE | DOI: 10.1109/ICRA48891.2023.10160988

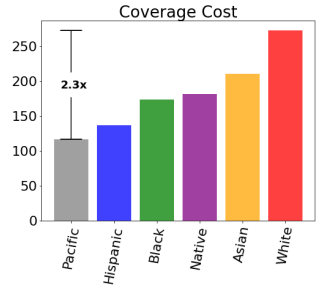
Authorized licensed use limited to: University of Pennsylvania. Downloaded on October 29,2026 at 14:38:57 UTC from IEEE Xplore. Restrictions apply.



(a) The population density ϕ in Philadelphia from U.S. census data, and positions of a 3 robot team using Lloyd's algorithm.



(b) The population densities ϕ_g by race as reported in the U.S. census, showing drastically different population sizes and spatial distributions among the different groups.



(c) The coverage cost per group based on the deployment of robots in 1a, showing up to a 2.3x difference in costs between groups.

Fig. 1: Population density functions are spatially complex. Additionally, 1b shows that the density functions of groups differ drastically, as do their population rates. 1c shows that Lloyd's algorithm can result in unfair outcomes: the worst off group has poor coverage, with a cost that is 2.3x higher than the cost of the best off group. Fair coverage would result in a lower worst-off group cost and less disparity in costs among the groups.

This formulation is common in literature including for societal applications like disaster response [4]–[7].

Consider a team of N robots $\mathcal{N} = \{1, \ldots, N\}$ whose positions and velocities at time t are $\mathbf{p}_{i,t} \in \mathbb{R}^2$ and $\mathbf{u}_{i,t} \in \mathcal{U} \subseteq \mathbb{R}^2$, respectively, with dynamics given by $\dot{\mathbf{p}}_{i,t} = \mathbf{u}_{i,t}, i \in \mathcal{N}$. The robots are equipped with sensors that can detect events; sensor performance degrades over distance [8], represented by $f(||\mathbf{q}-\mathbf{p}||) : \mathbb{R}^+ \to \mathbb{R}^+, \mathbf{q} \in \mathcal{D}$, which can also be interpreted as "the cost of moving to service an event" [4]. The function f is continuously differentiable, strictly positive, and strictly increasing [8]. Commonly $f(||\mathbf{q}-\mathbf{p}||) := ||\mathbf{q}-\mathbf{p}||^2$, which we use for evaluation in Section V.

The typical coverage control approach uses Voronoi iteration, also known as Lloyd's algorithm, which partitions the environment and then defines the control input for each robot as the gradient of a coverage cost function [4]–[7]. To calculate the Voronoi partitions, it is assumed that robots know the relative locations of their neighbors. To determine control inputs, each robot uses the density function within its partition; while some works investigate unknown ϕ, we assume the robot team has prior knowledge of the density function.

Voronoi cells are all points closest to each robot: $\mathcal{V}_{i,t} = \{\mathbf{q} \in \mathcal{D} : \|\mathbf{p}_{i,t} - \mathbf{q}\| \leq \|\mathbf{p}_{j,t} - \mathbf{q}\|, \forall j \in \mathcal{N}\}$. The coverage cost function measures the average sensing quality or cost to move, f, weighted by the density ϕ from each robot to all points in its Voronoi cell. For $f = ||\mathbf{q}-\mathbf{p}||^2$, the coverage cost function is interpreted as the weighted-average distance. The coverage cost is formally defined as:

$$\mathcal{J}_\phi\left(\mathbf{p}_{1,t} \ldots \mathbf{p}_{N,t}\right) = \sum_{i \in \mathcal{N}} \int_{\mathcal{V}_{i,t}} f\left(\|\mathbf{p}_{i,t} - \mathbf{q}\|\right) \phi(\mathbf{q}) \mathbf{dq} \quad (1)$$

Coverage control seek to minimize $\mathcal{J}_\phi(\mathbf{p}_T)$ at time horizon T. Therefore, Lloyd's algorithm defines a gradient descent law for the control inputs of the robots:

$$\dot{\mathbf{p}}_{i,t} = -\frac{\partial \mathcal{J}_\phi}{\partial \mathbf{p}_{i,t}} \quad (2)$$

This control law is guaranteed to converge to a local optimum where each robot is located at the weighted center of mass of its Voronoi cell, known as a *centroidal Voronoi tessellation* [9]. For the remainder of this paper, we will use "Lloyd's algorithm" to refer to this traditional approach.

B. Related Work

Fairness has been studied extensively in the artificial intelligence community [10], [11] with a strong focus on fairness in classification [12]–[14]. Definitions of group fairness such as demographic parity [11] and conditional statistical parity [15] leverage notions of independence between sensitive characteristics and predictions. Similarly, social fair clustering has been defined as achieving independence between sensitive characteristics and clustering cost [16], [17]. In this work, we similarly leverage the independence criterion of fairness, **minimizing the maximum coverage cost among a set of groups**, such as socio-economic, age, or racial groups.

Recently, the robotics community has begun investigating fairness, translating definitions of distributional fairness to multi robot applications such as task allocation [18], cooperative target tracking [19], and coverage path planning [20]. Note that coverage path planning seeks paths that visit all points in a region [21], as opposed to coverage control, where robots spread out to stationary points to monitor for events [22]. Lastly, researchers have highlighted the importance of fairness and ethics in robotics application like search and rescue [23]–[25]. To the best of our knowledge, we are the first to investigate fairness in coverage control.

Coverage control has been studied extensively, generally investigating one of three categories: (1) heterogeneity, considering differing robot sizes [26], mobility constraints [1], and sensing capabilities [7] ; (2) time-varying scenarios, where the density function ϕ is not static [27], [28]; and (3) learning based methods, where ϕ is unknown a priori [29], [30]. In this work, we consider a **homogeneous** team of robots **covering static, known density functions**. The dimensions of coverage control listed above add complexity to fairness and will be considerations of future work.

Lastly, for practitioners in coverage control, note that many definitions of fairness, including that presented in this paper, use min-max approaches. Mathematically, these fairness methods can be akin to providing robustness against worst

Authorized licensed use limited to: University of Pennsylvania. Downloaded on October 29,2026 at 14:38:57 UTC from IEEE Xplore. Restrictions apply.

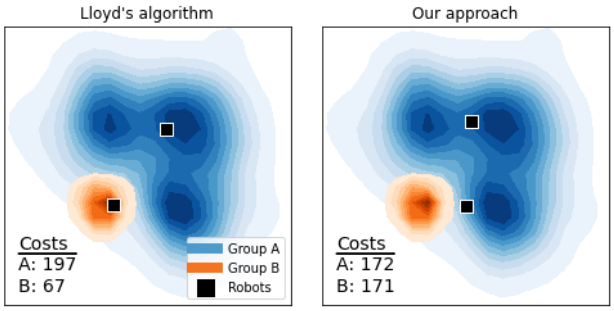


Fig. 2: The normalized density functions of two groups. (Left) The final coverage positions of two robots using Lloyd's algorithm. (Right) The positions resulting from our approach. While the traditional coverage approach achieves a better average, there is a *3x difference* in costs between the groups, despite normalizing for differing group population rates.

case outcomes [18]. While existing work has considered coverage for multiple event types [31], methods have not been presented in the literature that provide worst case robustness over a set of objectives. Our contributions not only investigate fairness, but also provide approaches for studying robust coverage control.

III. Problem Formulation

Consider coverage control applications like Example 1 where the events described by the density function ϕ are people in a social context. Therefore, these events intrinsically have the same sensitive characteristics that describe the population (e.g., race, age, sex, religion). Given a set $\mathcal{G}$ of possible groups, every possible event belongs to one group $g \in \mathcal{G}$.[1] While ϕ describes the likelihood of any event, ϕ_g is the likelihood of an event from group g.

Our goal is to develop coverage control algorithms that provide fair coverage with respect to sensitive characteristics in the population, namely, that the coverage cost is independent of the group label g. The remainder of this section investigates the disparate impact of the existing methods described in Section II-A and then presents a novel problem formulation that is the first to define fairness in coverage control.

A. Disparate Impact of Existing Methods

To investigate the coverage performance of groups within a population, we use Lloyd's algorithm for coverage control (as described in Section II-A) on real population data for Philadelphia County, Pennsylvania, USA (details on the data and simulations in Section V). The density function ϕ is defined as the population density, as seen in Figure 1a. The group density functions ϕ_g are defined using census data on population rates of racial groups, as seen in Figure 1b. Lloyd's algorithm is run until convergence using ϕ. Then, the coverage cost is calculated for ϕ and for all groups ϕ_g, $\forall g \in \mathcal{G}$. Figure 1 provides a clear example of the disparate costs that result from the classical coverage control approach. Figure 1c shows that coverage costs vary drastically among the groups, with a $2.3x$ difference between the best-off and worst-off groups.

Two factors that impact disparate outcomes are differing *population rates* and differing *spatial distributions* among the groups. Since the density functions are known a priori, *normalized* group density functions could compensate for differing population rates, i.e., $\phi = \frac{1}{|\mathcal{G}|}\sum_{g\in\mathcal{G}}\phi_g$ where each ϕ_g has a total probability mass of 1. However, the clustering literature shows that, for the discrete Lloyd's algorithm, reweighting to compensate for group rates does not achieve fairness [16], [17]. Additionally, our counterexample in Figure 2 shows that normalization does not achieve fairness; the resulting group costs of Lloyd's algorithm differ by a factor of 3. Our approach yields nearly equal cost for the two groups.

Even though ϕ can be designed to cater to different scenarios or applications, it cannot be designed to capture fairness with respect to sensitive characteristics in a population. Therefore, we define a new objective function in place of Equation (1) to capture the notion of fair coverage cost.

B. Socially Fair Coverage Problem Definition

Many definitions of algorithmic fairness are based on a notion of independence between an algorithm's outputs and the sensitive characteristics of the inputs [10]. As such, our goal is to develop coverage control algorithms whose coverage cost is independent of sensitive characteristics. Ultimately, we strive to provide more equitable coverage to groups while maintaining usefulness of the coverage control.

Optimizing strictly for equal coverage cost among the groups can yield poor performance for all groups; trivially, robots infinitely far away from the population produces equitable, though arbitrarily poor, outcomes for all groups. Instead, we use a **min-max approach, optimizing the worst-off group's coverage cost**. This objective drives the system towards improved coverage for all groups while reducing disparate outcomes by focusing on improving the coverage for the group with the highest coverage cost.

Problem 1 (Socially Fair Coverage)**.** *Given a spatial region $\mathcal{D} \subset \mathbb{R}^2$ and a set of density functions corresponding to the set of groups within a population $\phi_g : \mathcal{D} \to \mathbb{R}^+, \forall g \in \mathcal{G}$, distribute a team of N robots in order to minimize the coverage cost of the group g with the worst coverage quality. Formally, Socially Fair Coverage* ***minimizes*** *the following objective:*

$$\mathcal{J} = \max_{g\in\mathcal{G}} \sum_{i\in\mathcal{N}} \int_{\mathcal{V}_{i,t}} f\left(\|\mathbf{p}_{i,t} - \mathbf{q}\|\right)\phi_g(\mathbf{q})\mathbf{dq} \qquad (3)$$

IV. Approach

Similar to the methods in Section II-A, we use Voronoi iteration to solve Problem 1. This approach requires a gradient descent law for defining the control inputs of the robots in order to minimize the coverage objective. However, the objective function $\mathcal{J}$ in Equation (3) does not readily afford a gradient based controller because the **maximum function is non-differentiable**. The remainder of this section

[1]It is common in the fairness literature to assume each event can only belong to one group $g \in \mathcal{G}$ [16], [17].

Authorized licensed use limited to: University of Pennsylvania. Downloaded on October 29,2026 at 14:38:57 UTC from IEEE Xplore. Restrictions apply.

approximates Equation (3), defines a gradient descent control law, and examines controller properties.

A. Gradient Controller of Approximate Objective Function

We define a gradient descent law for Voronoi iteration by using a smooth approximation of the objective function $\mathcal{J}$ in Equation (3). By combining Equations (1) and (3), we see that $\mathcal{J} = \max_{g\in\mathcal{G}} \mathcal{J}_{\phi_g}$. Similar maximum function are often approximated in the machine learning literature by the **LogSumExp (LSE)** [32]–[34], yielding the smooth function:

$$\mathcal{J} \approx \hat{\mathcal{J}} = \log \sum_{g\in\mathcal{G}} \exp \mathcal{J}_{\phi_g} \tag{4}$$

The partial derivative of $\hat{\mathcal{J}}$ with respect to a robot's location $\mathbf{p}_{i,t}$, is the weighted sum of the partial derivatives of $\mathcal{J}_{\phi_g}$ with respect to the robot's location:

$$\frac{\partial \hat{\mathcal{J}}}{\partial \mathbf{p}_{i,t}} = \sum_{g\in\mathcal{G}} \delta_g \frac{\partial \mathcal{J}_{\phi_g}}{\partial \mathbf{p}_{i,t}}, \quad \delta_g = \frac{\exp \mathcal{J}_{\phi_g}}{\sum_{g\in\mathcal{G}} \exp \mathcal{J}_{\phi_g}} \tag{5}$$

Details on the derivation of Equation (5) can be found in Appendix A. To solve Problem 1 using Voronoi iteration, we define the following a gradient based controller:

$$\dot{\mathbf{p}}_{i,t} = -\frac{\partial \hat{\mathcal{J}}}{\partial \mathbf{p}_{i,t}} \tag{6}$$

B. Controller Properties

We examine the impact of the approximation made in Equation (4) on both $\mathcal{J}$ and the respective gradient based controller, the global information needed for our approach, and the convergence properties of the controller.

Approximation. The LogSumExp (LSE) over approximates the maximum function with a known bound of $\log|\mathcal{G}|$. Because this error bound is fixed, the *relative* error is greater for small values of $\mathcal{J}_{\phi_g}$. However, the functions can be scaled to achieve tighter bounds: $\frac{1}{M}\log\sum_{g\in\mathcal{G}}\exp M\mathcal{J}_{\phi_g}$ has a bound of $\frac{\log|\mathcal{G}|}{M}$.

The approximation of $\mathcal{J}$ impacts the controller through the values of δ_g. For a group g_0 such that $\mathcal{J}_{\phi_{g_0}} \ll \sum_{g\in\mathcal{G}\setminus g_0} \mathcal{J}_{\phi_g}$, the value of δ_{g_0} tends toward zero. When $\mathcal{J}_{\phi_{g_0}} \gg \sum_{g\in\mathcal{G}\setminus g_0} \mathcal{J}_{\phi_g}$, the value of δ_{g_0} tends toward one. This behavior aligns with the expected output for a hypothetical controller of the exact objective $\mathcal{J}$; the δ_g weights would be binary values indicating whether a group had the worst coverage cost.

We see that Equation (6) derived from our approximation of $\mathcal{J}$ deviates from the hypothetical exact controller when $\mathcal{J}_{\phi_g}$ values are similar. E.g., when all $\mathcal{J}_{\phi_g}$ values are the same, $\delta_g = \frac{1}{|\mathcal{G}|}\forall g \in \mathcal{G}$. Therefore, similarly valued $\mathcal{J}_{\phi_g}$ yield a controller that improves the average coverage cost across the groups. Despite deviating from the strict min-max definition of fairness in Problem 1, this behavior is beneficial: the controller improves fairness when groups have highly disparate costs, and when group costs are similar, the controller prevents instability and optimizes average overall performance.

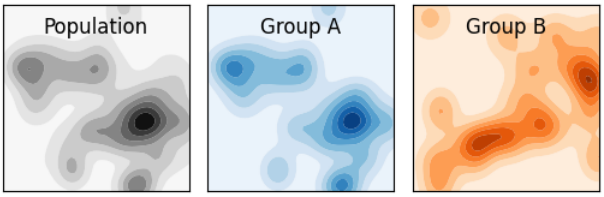


Fig. 3: An example of the random Gaussian Mixture Model environments where the ratio between group A and B is 3:1, meaning the minority is 25% of the population.

Global Information. Traditional Voronoi iteration approaches to the coverage control problem are distributed; the control input for a robot i at any time t is calculated using only the information in that robot's partition $\mathcal{V}_{i,t}$, i.e., $\frac{\partial \mathcal{J}_\phi}{\partial \mathbf{p}_{i,t}}$ depends only on the locations of a robots neighboring teammates.

In our approach, the control input with respect to a group, $\frac{\partial \mathcal{J}_{\phi_g}}{\partial \mathbf{p}_{i,t}}$, depends only on neighbor locations and can therefore be calculated in a distributed manner. However, the entire control input $\frac{\partial \hat{\mathcal{J}}}{\partial \mathbf{p}_{i,t}}$ does not depend only on neighbor locations; the weights δ_g depend on the global coverage cost of all groups which is a function of the location of all N robots on the team.

Convergence. An important property of Voronoi iteration is that the robots converge to a stationary configuration. We show that the control law induced by Equation (6) converges.

$$\frac{d\hat{\mathcal{J}}}{dt} = \sum_{i\in\mathcal{N}} \frac{\partial \hat{\mathcal{J}}}{\partial \mathbf{p}_i} \dot{\mathbf{p}}_i^T = -\sum_{i\in\mathcal{N}} \left\| \frac{\partial \hat{\mathcal{J}}}{\partial \mathbf{p}_i} \right\|^2 \tag{7}$$

The time derivative of $\hat{\mathcal{J}}$ is the negative sum over all robots of $\|\frac{\partial \hat{\mathcal{J}}}{\partial \mathbf{p}_i}\|$ squared. Therefore, the derivative in Equation (7) is non-positive. Also, $\frac{d\hat{\mathcal{J}}}{dt} = 0$ if and only if $\|\frac{\partial \hat{\mathcal{J}}}{\partial \mathbf{p}_i}\| = 0\ \forall i \in N$. Therefore, $\hat{J}$ monotonically decreases until all robots reach a stationary point. Therefore, the control law induced by Equation (6) converges to a local minimum as $t \to \infty$.

V. Evaluation

We create two simulated environments: random Gaussian mixture models, and kernel-density estimates of U.S. census data. The first studies performance across differing group sizes; larger relative size between groups increases the group disparity for traditional methods but does not impact our approach. The latter set of experiments validates our approach on real data; our method generally decreases disparity of coverage costs between groups compared to Lloyds while maintaining a similar computation and convergence time.

A. Randomized Environments

We use Gaussian mixture models (GMMs) to generate random environments to test our methods against traditional coverage control (referred to as "Lloyds"). Each environment has two groups represented by a finite number of Gaussian distributions. We vary the relative number of distributions between the two groups to create different majority:minority ratios. We test ratios ranging from minority group sizes of

Authorized licensed use limited to: University of Pennsylvania. Downloaded on October 29,2026 at 14:38:57 UTC from IEEE Xplore. Restrictions apply.

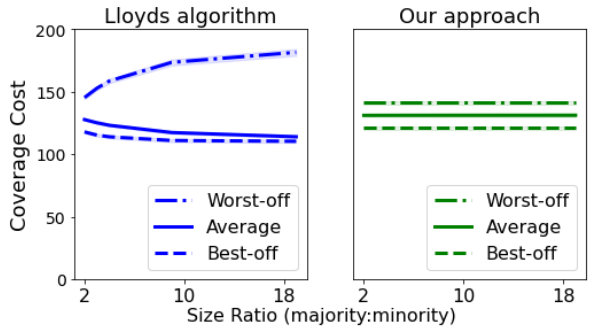


Fig. 4: The average and group coverage costs as a function of the majority to minority size ratio (mean and 95% confidence). As the size difference between the groups grow, the disparate impact of Lloyd's algorithm grows. Our method achieves less disparate costs independent of size ratio.

33% down to 5%. To generate the GMMs, we choose random means for each group (the number of means corresponding to the majority:minority ratio). For each mean, we generate a random covariance. An example environment is seen in Figure 3.

We generate 500 random environments and run Lloyds and our approach until convergence for a team of 5 robots. We report the average cost and the cost per group as a function of the ratio of group sizes in Figure 4. We found that larger differences in the size of the groups leads to larger disparate coverage cost when using the traditional method. Additionally, Lloyds algorithm yields a decreasing average cost as the size ratio between groups increases **at the expense of the minority group**: the average cost and the majority group cost decrease slightly while the minority group cost increases drastically. In contrast, our method provides lower disparity between the groups, regardless of the size difference of the groups.

B. Census Data Experiments

We validate our method using data from *census.gov* on population and demographics for U.S. counties. We generate test environments by defining a convex spatial region $\mathcal{D}$ that encompasses county shapes.[2][3] To generate the density functions ϕ and ϕ_g, we use demographics and population data.[4] The density functions ϕ_A and ϕ_B represent the white and non-white[5] populations, respectively. Census data is reported per census tract (a small subdivisions of counties). Therefore, we uniformly sample census tracts to estimate the population across the entire county. We use this set of samples to create a kernel density estimate, yielding a smooth density function.

We generate environments of the 20 most populous counties in the country, omitting non-continuous counties. For each county, we implement Lloyd's algorithm and our approach for $N = 3$ robots. We run the algorithms until convergence and measure the groups costs, total population cost, number steps to converge, and total computation time.

[2]*census.gov/geographies/mapping-files/time-series/geo/tiger-line-file.html*

[3]Non-compact counties are ommitted.

[4]*census.gov/programs-surveys/acs/data.html*

[5]Non-white includes all races reported by the census other than white.

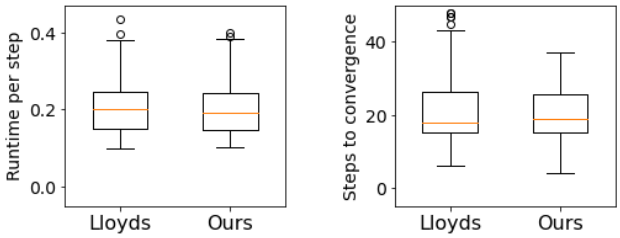


Fig. 5: The distribution of computation time and number of steps until convergence shows that our method has similar runtime and convergence performance as the traditional Lloyd's approaches.

Because the Voronoi iteration converges to local minima, we conduct 5 trials with different random starting locations. Reported values are aggregated means over the 5 trials.

For each of the 20 counties, Figure 6 reports three costs per approach: the worst-off group cost, population wide cost, and best-off group cost. The worst and best-off costs are $\max\{\mathcal{J}_{\phi_A}, \mathcal{J}_{\phi_B}\}$ and $\min\{\mathcal{J}_{\phi_A}, \mathcal{J}_{\phi_B}\}$ because neither group received worse costs for all counties. Our approach is group agnostic, optimizing for whichever group is worst off. The average population cost $\mathcal{J}_{\phi_A}$ is not equivalent to the average of $\mathcal{J}_{\phi_A}$ and $\mathcal{J}_{\phi_B}$ because the size of A and B differ.

Our approach outperforms Lloyds with respect to fairness for the majority of counties tested: the worst-off group's cost is lower as is the cost disparity between groups. Notably, some counties like *San Bernardino County* and *Harris County* have a large disparity in group costs. In both cases, our approach significantly reduces the worst-off group cost. In the former, our approach yields slightly worse average cost compared to Lloyds. However, in the latter, our approach achieves a better average cost in addition to the improved maximum cost.

Figure 6 also shows cases where our approach and Lloyd's algorithm produce similar costs, and our approach leads to slightly worse outcomes. This negative result is likely because Voronoi iteration (which is used in both approaches) is a gradient based solution to an NP-hard problem which converges to a local minimum. Further investigation is needed to determine whether aspects of the spatial distributions are correlated with the different levels of performance (Section VI). Lastly, we report the number of steps until convergences[6] and the average computation time per step. In Figure 5, we see that our approach has a similar convergence time and computation time compared to the traditional approach.[7] Therefore, we achieve increased fairness without additional computational cost.

VI. Discussion

Scenarios like Example 1 demonstrate the need to consider fairness. Conversely, for applications like scientific exploration, fairness may seem irrelevant. However, these scenarios often have multiple stakeholders, such as scientific teams of different disciplines conducting a joint oceanic exploration

[6]Experiments are implemented in a discrete time simulation.

[7]Omitted trials that reached the simulation limit, with robots bouncing around a minimum (less than 5% of trials for both approaches).

Authorized licensed use limited to: University of Pennsylvania. Downloaded on October 29,2026 at 14:38:57 UTC from IEEE Xplore. Restrictions apply.

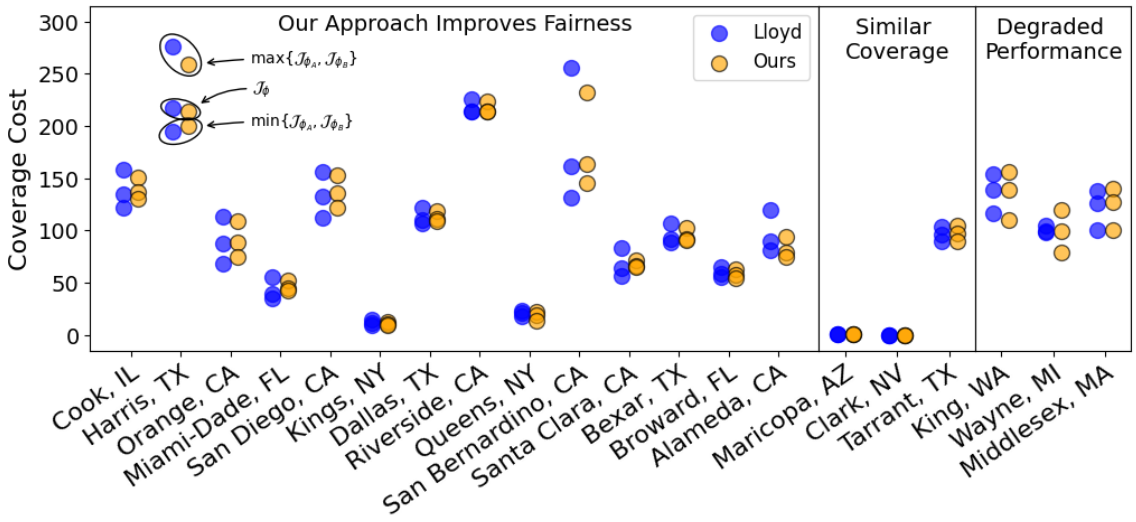


Fig. 6: Evaluation of Lloyd's algorithm (blue) and our fair approach (orange) for the 20 most populous counties in U.S. For each city and approach, the three plotted points represent: the worst-off group cost $\max\{\mathcal{J}_{\phi_A}, \mathcal{J}_{\phi_B}\}$, the population cost $\mathcal{J}_\phi$, and the best-off group cost $\min\{\mathcal{J}_{\phi_A}, \mathcal{J}_{\phi_B}\}$. For most counties, our approach reduces the worst-off cost and reduces cost disparity between the groups, with little impact on the overall population cost $\mathcal{J}_\phi$.

mission. The different needs of these stakeholders can be modeled by a set of importance functions $\phi_g, \ \forall g \in \mathcal{G}$. Therefore, the methods we presented above can be used to balance the needs of these stakeholders in a fair manner.

Robustness. The methods developed in this work can be used to provide worst case robustness over a set of objectives. This type of robustness has not yet been studied in the coverage control literature. While there is some mathematical overlap between the methods of fairness and robustness, further research is needed on both fronts. For example, population and demographics data has embedded in it biases from policies such as redlining which may not apply in situations studying robustness. Further investigation is needed to discern the causes of unfairness in coverage control beyond differing population rates and spatial distributions. Integral to the future work of fairness in coverage control is the investigation of different fairness definitions for these contexts.

Approximation. Future work will develop new methods for solving Problem 1. First, approximations other than LogSumExp or gradient based methods for non-differentiable function may be developed. Also, future work can develop controllers for heterogeneous teams, such as for robots with heterogeneous mobility constraints. Lastly, many applications, like disaster response, have dynamic or unknown density functions. This partial observability adds additional complexity to considerations of fairness and to the design of controllers.

VII. Conclusion

We show that existing coverage control methods can cause disparate coverage costs with respect to sensitive characteristics in a population in societal applications such as disaster response. Therefore, fairness is an important consideration that the coverage control community must investigate.

In this work, we propose the first definition of social fairness for coverage control as the minimization of the maximum coverage cost over a set of density functions and we solve for this objective using Voronoi iteration. To derive a controller, we approximate this non-differentiable objective function using the LogSumExp. Our resulting controller is the weighted sum of control inputs with respect to each group; the weights depend on the relative global coverage cost of each group. When group costs are highly disparate, our approach decrease the cost of the worst-off groups, improving fairness. When group costs are similar, our approach decreases the average cost amongst the group, thus improving efficiency.

We demonstrate our approach in simulated environments of randomly generated GMMs and kernel-density estimates of U.S. census data. We show that our method: (1) retains similar computational cost and convergence properties as existing methods, (2) provides fairness irrespective of disparate population rates among groups, and (3) improves fairness on real population data.

Appendix

The following is the derivation of Equation (5), which uses the derivative of the log function $\frac{\partial}{\partial x}\log[f(x)] = \frac{1}{f(x)}f'(x)$ and of the exponential function $\frac{\partial}{\partial x}\exp[f(x)] = \exp[f(x)]f'(x)$

$$\hat{\mathcal{J}} = \log \sum_{g\in\mathcal{G}} \exp \mathcal{J}_{\phi_g}$$

$$\frac{\partial \hat{\mathcal{J}}}{\partial \mathbf{p}_{i,t}} = \frac{1}{\sum_{g\in\mathcal{G}} \exp \mathcal{J}_{\phi_g}} \frac{\partial}{\partial \mathbf{p}_{i,t}} \left(\sum_{g\in\mathcal{G}} \exp \mathcal{J}_{\phi_g} \right)$$

$$= \frac{1}{\sum_{g\in\mathcal{G}} \exp \mathcal{J}_{\phi_g}} \left(\sum_{g\in\mathcal{G}} \exp \mathcal{J}_{\phi_g} \frac{\partial \mathcal{J}_{\phi_g}}{\partial \mathbf{p}_{i,t}} \right)$$

$$= \sum_{g\in\mathcal{G}} \delta_g \frac{\partial \mathcal{J}_{\phi_g}}{\partial \mathbf{p}_{i,t}}, \quad \delta_g = \frac{\exp \mathcal{J}_{\phi_g}}{\sum_{g\in\mathcal{G}} \exp \mathcal{J}_{\phi_g}}$$

Authorized licensed use limited to: University of Pennsylvania. Downloaded on October 29,2026 at 14:38:57 UTC from IEEE Xplore. Restrictions apply.

REFERENCES

[1] Y. Shi, N. Wang, J. Zheng, Y. Zhang, S. Yi, W. Luo, and K. Sycara, "Adaptive Informative Sampling with Environment Partitioning for Heterogeneous Multi-Robot Systems," *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2020.

[2] W. Gosrich, S. Mayya, R. Li, J. Paulos, M. Yim, A. Ribeiro, and V. Kumar, "Coverage control in multi-robot systems via graph neural networks," in *2022 International Conference on Robotics and Automation (ICRA)*. IEEE, 2022, pp. 8787–8793.

[3] J. Miller and J. P. How, "Predictive positioning and quality of service ridesharing for campus mobility on demand systems," in *2017 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2017, pp. 1402–1408.

[4] J. Cortes, "Coverage Optimization and Spatial Load Balancing by Robotic Sensor Networks," *IEEE Transactions on Automatic Control*, 2010.

[5] S. G. Lee, Y. Diaz-Mercado, and M. Egerstedt, "Multirobot control using time-varying density functions," *IEEE Transactions on robotics*, vol. 31, no. 2, pp. 489–493, 2015.

[6] J. Cortes, S. Martinez, T. Karatas, and F. Bullo, "Coverage control for mobile sensing networks," *IEEE Transactions on Robotics and Automation*, 2002.

[7] M. Santos, Y. Diaz-Mercado, and M. Egerstedt, "Coverage Control for Multirobot Teams With Heterogeneous Sensing Capabilities," *IEEE Robotics and Automation Letters*, vol. 3, no. 2, pp. 919–925, Apr. 2018.

[8] K. Laventall and J. Cortes, "Coverage control by multi-robot networks with limited-range anisotropic sensory," *International Journal of Control*, 2009.

[9] Q. Du, V. Faber, and M. Gunzburger, "Centroidal voronoi tessellations: Applications and algorithms," *SIAM review*, vol. 41, no. 4, pp. 637–676, 1999.

[10] N. Mehrabi, F. Morstatter, N. Saxena, K. Lerman, and A. Galstyan, "A survey on bias and fairness in machine learning," *ACM Computing Surveys (CSUR)*, vol. 54, no. 6, pp. 1–35, 2021.

[11] C. Dwork, M. Hardt, T. Pitassi, O. Reingold, and R. Zemel, "Fairness through awareness," in *Proceedings of the 3rd innovations in theoretical computer science conference*, 2012, pp. 214–226.

[12] M. Kearns, S. Neel, A. Roth, and Z. S. Wu, "Preventing fairness gerrymandering: Auditing and learning for subgroup fairness," in *International Conference on Machine Learning*. PMLR, 2018, pp. 2564–2572.

[13] S. Verma and J. Rubin, "Fairness definitions explained," in *2018 ieee/acm international workshop on software fairness (fairware)*. IEEE, 2018, pp. 1–7.

[14] A. K. Menon and R. C. Williamson, "The cost of fairness in binary classification," in *Conference on Fairness, Accountability and Transparency*. PMLR, 2018, pp. 107–118.

[15] S. Corbett-Davies, E. Pierson, A. Feller, S. Goel, and A. Huq, "Algorithmic decision making and the cost of fairness," in *Proceedings of the 23rd acm sigkdd international conference on knowledge discovery and data mining*, 2017, pp. 797–806.

[16] M. Ghadiri, S. Samadi, and S. Vempala, "Socially Fair k-Means Clustering," *arXiv: Learning*, 2020.

[17] M. Abbasi, A. Bhaskara, and S. Venkatasubramanian, "Fair clustering via equitable group representations," in *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, 2021, pp. 504–514.

[18] M. Malencia, V. Kumar, G. Pappas, and A. Prorok, "Fair robust assignment using redundancy," *IEEE Robotics and Automation Letters*, vol. 6, no. 2, pp. 4217–4224, 2021.

[19] J. Banfi, J. Guzzi, A. Giusti, L. Gambardella, and G. A. Di Caro, "Fair multi-target tracking in cooperative multi-robot systems," in *2015 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2015, pp. 5411–5418.

[20] M. Brandao, "Socially Fair Coverage: The Fairness Problem in Coverage Planning and a New Anytime-Fair Method," *2021 IEEE International Conference on Advanced Robotics and Its Social Impacts (ARSO)*, 2021.

[21] E. Galceran and M. Carreras, "A survey on coverage path planning for robotics," *Robotics and Autonomous systems*, vol. 61, no. 12, pp. 1258–1276, 2013.

[22] J. Cortes, S. Martinez, and F. Bullo, "Spatially-distributed coverage optimization and control with limited-range interactions," *arXiv: Optimization and Control*, 2004.

[23] L. Battistuzzi, C. T. Recchiuto, and A. Sgorbissa, "Ethical concerns in rescue robotics: a scoping review," *Ethics and Information Technology*, vol. 23, no. 4, pp. 863–875, 2021.

[24] M. Brandao, M. Jirotka, H. Webb, and P. Luff, "Fair navigation planning: A resource for characterizing and designing fairness in mobile robots," *Artificial Intelligence*, vol. 282, p. 103259, 2020.

[25] J. V. Hurtado, L. Londoño, and A. Valada, "From learning to relearning: A framework for diminishing bias in social robot navigation," *Frontiers in Robotics and AI*, vol. 8, p. 650325, 2021.

[26] O. Arslan and D. E. Koditschek, "Voronoi-based coverage control of heterogeneous disk-shaped robots," in *2016 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2016, pp. 4259–4266.

[27] S. G. Lee and M. Egerstedt, "Controlled Coverage Using Time-Varying Density Functions," *IFAC Proceedings Volumes*, 2013.

[28] M. Santos, S. Mayya, G. Notomista, and M. Egerstedt, "Decentralized Minimum-Energy Coverage Control for Time-Varying Density Functions," *2019 International Symposium on Multi-Robot and Multi-Agent Systems (MRS)*, 2019.

[29] W. Luo, C. Nam, G. Kantor, and K. Sycara, "Distributed Environmental Modeling and Adaptive Sampling for Multi-Robot Sensor Coverage," *AAMAS*, 2019.

[30] L. Wei, A. McDonald, V. Srivastava, and V. Srivastava, "Multi-Robot Gaussian Process Estimation and Coverage: Deterministic Sequencing Algorithm and Regret Analysis," *2021 IEEE International Conference on Robotics and Automation (ICRA)*, 2021.

[31] A. Sadeghi and S. L. Smith, "Coverage Control for Multiple Event Types with Heterogeneous Robots," in *2019 International Conference on Robotics and Automation (ICRA)*. Montreal, QC, Canada: IEEE, May 2019, pp. 3377–3383. [Online]. Available: https://ieeexplore.ieee.org/document/8793639/

[32] F. Nielsen and K. Sun, "Guaranteed bounds on the kullback–leibler divergence of univariate mixtures," *IEEE Signal Processing Letters*, vol. 23, no. 11, pp. 1543–1546, 2016.

[33] C. Shen and H. Li, "On the dual formulation of boosting algorithms," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 32, no. 12, pp. 2216–2231, 2010.

[34] P. Blanchard, D. J. Higham, and N. J. Higham, "Accurately computing the log-sum-exp and softmax functions," *IMA Journal of Numerical Analysis*, vol. 41, no. 4, pp. 2311–2330, 2021.

Authorized licensed use limited to: University of Pennsylvania. Downloaded on October 29,2026 at 14:38:57 UTC from IEEE Xplore. Restrictions apply.