Socially Fair Coverage: The Fairness Problem in Coverage Planning and a New Anytime-Fair Method

Martim Brandão

***Abstract*— In this paper we investigate and characterize social fairness in the context of coverage path planning. Inspired by recent work on the fairness of goal-directed planning, and work characterizing the disparate impact of various AI algorithms, here we simulate the deployment of coverage robots to anticipate issues of fairness. We show that classical coverage algorithms, especially those that try to minimize average waiting times, will have biases related to the spatial segregation of social groups. We discuss implications in the context of disaster response, and provide a new coverage planning algorithm that minimizes cumulative unfairness at all points in time. We show that our algorithm is 200 times faster to compute than existing evolutionary algorithms—while obtaining overall-faster coverage and a fair response in terms of waiting-time and coverage-pace differences across multiple social groups.**

I. Introduction

Coverage planning is an important problem with applications in inspection, surveillance, disaster response and other domains. Many of these applications are socially charged: for example, search missions in disaster response need to cover whole impacted areas but they also need to respect aspects of priority and fairness [1], [2]. Disaster response missions are also expected to attend to the most-at-risk first [2], [3], [4], and not to penalize already-marginalized communities [5]. Similarly, wildfire tracking [6], surveillance [7] and security [8] can also affect different communities differently and thus lead to concerns of fairness.

In this paper we investigate whether existing coverage planning algorithms can raise concerns of fairness, and propose a new method to alleviate such issues. We are inspired by recent work [9] which demonstrated the existence of a fairness dimension to goal-directed path planning—particularly its capacity to reinforce spatially-correlated social inequalities. However, the analysis in [9] does not translate directly to the coverage problem since by definition coverage paths will serve the whole population of interest. As we will show in this paper, what matters in coverage planning are relative waiting times rather than relative final coverage of different social groups. Importantly, we show that coverage paths produced by classical algorithms can lead to a bias in waiting times. This penalizes spatially segregated social groups such as ethnic minorities and age-groups living in lowly-populated areas. Such issues could be fatal to the acceptance of AI-assisted drone technology in disaster response, or further fuel existing criticism regarding the impact that responses have on already-marginalized groups [5].

*Martim Brandão is with King's College London, London, UK.

Based on this investigation, in this paper we provide a concept of fairness relevant to the coverage planning problem—"anytime-fairness"—and we propose a new method that handles the dimensionality of the problem better than state-of-the-art evolutionary methods.

Our contributions are the following:

1) We use realistic simulations of robot deployment to characterize issues of fairness in disaster response coverage-based robots;
2) We propose a new fairness-aware method for coverage path planning that is both computationally faster than state-of-the-art methods, coverage-efficient, and leads to a balanced distribution of waiting times across social groups despite spatial segregations.

II. Related work

Coverage path planning is the problem of computing a path over a graph that visits all nodes. Many methods have been proposed to tackle such problem, from Boustrophedon's back-and-forth ox-like paths [10] to more recent evolutionary algorithms [7]. We refer the interested reader to recent surveys on coverage algorithms [11].

This paper is also related to the work of Brandão et al. [9], which proposes a new Responsible Innovation methodology to identify issues of fairness by simulating robot deployment. Here we take a similar approach to anticipating issues of fairness: we simulate the deployment of robots using coverage path planning algorithms, and analyze the relative impact to different social groups. Compared to [9] we focus on coverage planning rather than goal-directed planning. While [9] defines fairness as the match between the social make-up of a region of interest and the people found along a robot's path, such definition is not applicable to coverage planning—since coverage paths will by definition cover whole populations. Instead, here we focus on differences in the *order* or *speed* at which social groups are covered. Additionally, our planning method is based on sampling and a greedy strategy, which as we will show leads to much faster computation times and better coverage than a pure extension of the methods in [9].

Within the loosely-defined planning literature, the concept of fairness has been applied to Linear Temporal Logic in socially-agnostic ways [12], [13] (where it refers to serving multiple regions infinitely often), to vehicle routing workloads [14], to logistics problems of train and fleet scheduling [15], and to doctor scheduling problems [16] to satisfy the preferences of individual workers. Contrary to the technical work on planning above, fairness as studied in humanistic works of disaster response [3], [4], urban planning [17], environmental justice [18], and fairness-centered work in machine learning [19], [20], [21] is often concerned less with personal preferences and more with disparate impact across socio-economical groups. These groups could be related to gender, race, poverty, or other attributes relevant to discrimination and marginalization. In this paper our approach is consistent with the latter notions of fairness at the level of social groups. We specifically characterize the potential disparate impact across age and ethnicity groups, though our method is general and applicable to other categorical groupings.

While most recent work on fairness in Artificial Intelligence has focused on machine learning and big data [20], [21], [22], here we use such work as inspiration to analyze the potential issues of disparate impact in the context of robot planning algorithms in deployment. Similarly to ML-centered work, we focus on concerns of social fairness, discrimination and marginalization, although our datasets are in the form of government-collected census data [23], and our disparate impact an emerging feature of socio-economic factors behind the spatial organization of people [18]. Our work is similar in motivation to recent research on fairness within robot navigation [9], [24], though here we focus on the specific problem of coverage planning.

III. The problem: social bias in coverage plans

In this section we seek to answer the Research Question: can rescue coverage paths reinforce social inequalities present in the spacial organization of cities? This will serve to inform search and rescue teams deploying search-drones, as well as to motivate our own coverage algorithm focused on reducing such bias.

To answer our question we use the Responsible Innovation methodology proposed by [9]: we use robot deployment simulations together with real social data to predict issues of inequality and unfairness. In particular we simulate the deployment of a victim-search drone using coverage path planning algorithms on a real map. We use openly-available census data from the Office for National Statistics in England, which includes maps along with spatial distributions of population density, age, ethnicity and other variables [23]. The analysis shown here is for the city of Oxford, UK, for the sake of example—though similar aspects of spatial segregation of populations related to age, ethnicity, and socio-economic factors are widespread [18], [17]. The size of the Oxford map is of 300x300 cells, and each cell is annotated with total population density, as well as population per age and ethnicity bin. We assume a drone flying at 50m height, at a constant speed of 7m/s.

A. Simplistic unconstrained coverage

We first simulate the deployment of robots in the simplistic setting of coverage without battery limits. In this case algorithms such as Boustrophedon decomposition [10] or spiral-shaped coverage can be used. Boustrophedon decomposes a map into cells which are then covered with back-and-forth (ox-like) paths. Spiral paths simply proceed from the center of the map and outwards in spiral form, until the whole map is covered. For Boustrophedon decomposition we use an open-source implementation[1] [25] starting the path from the center of the map—the location of a fire station as the hypothetical drone-launching site.

Fig. 1 shows the coverage paths, cumulative population covered (according to the population density data), and then indicators of social inequality: per-social-group cumulative population coverage at particular times (20% of full coverage), and the range of these values throughout time (i.e. maximum minus minimum of group coverage at each time). Red lines on the maps are the paths taken by the hypothetical robot, green-colored areas are cells sensed by the robot, assuming the robot flies at 50m height with a 90degree downward-facing camera. Black areas are hence not sensed by the robot at any time. The city of Oxford has a zero-population area in the North-West, which is avoided by the Boustrophedon path, hence leaving an uncovered area in the map. The figure shows that the spiral-shaped path finds more people early on, which is due to high population density in the center of the city, though the Boustrophedon path eventually covers the full population earlier due to more efficient planning (fewer paths cross zero-population areas and already-visited map areas).

Interestingly, the spatial organization of the population is such that the younger student population (18-24 years) is concentrated in the center, near university buildings, while the older population is scattered throughout the rest of the city. This leads spiral paths to find roughly 70% of the undergraduate student population at the first one-fifth of the map, while infant and elderly populations are only 20% covered at this time. This corresponds to the 50% (70%-20%) visit range peak on the “Group coverage range” graph. Such inequality could potentially be of concern to rescue teams, since the elderly and infant populations are typically of higher risk—but they are only found later on in the mission, thus decreasing potential survival rates. The Boustrophedon path leads to a similarly high inequality peak (35%), since it also begins at the center of the city. The figure also shows similar inequality in terms of ethnicity: ethnic minorities such as Black, Indian and Arab are penalized by the coverage paths. The reason for this is again that the population in the center is mostly White or Chinese, while racial minorities such as Black, Arab and Indian populations are scattered outside the city center in racially segregated ways due to various historic and economic reasons. Again, leaving coverage of such groups towards the end could be concerning to rescue teams. Recent disaster response missions have been criticized for penalizing racial minorities and poorer populations, who are from the start less likely to survive or escape disaster areas on their own due to lower economic resources [5]. These simulations thus show that inherent spatial distributions of social groups can lead to a

[1] https://github.com/18alantom/CoveragePathPlanning

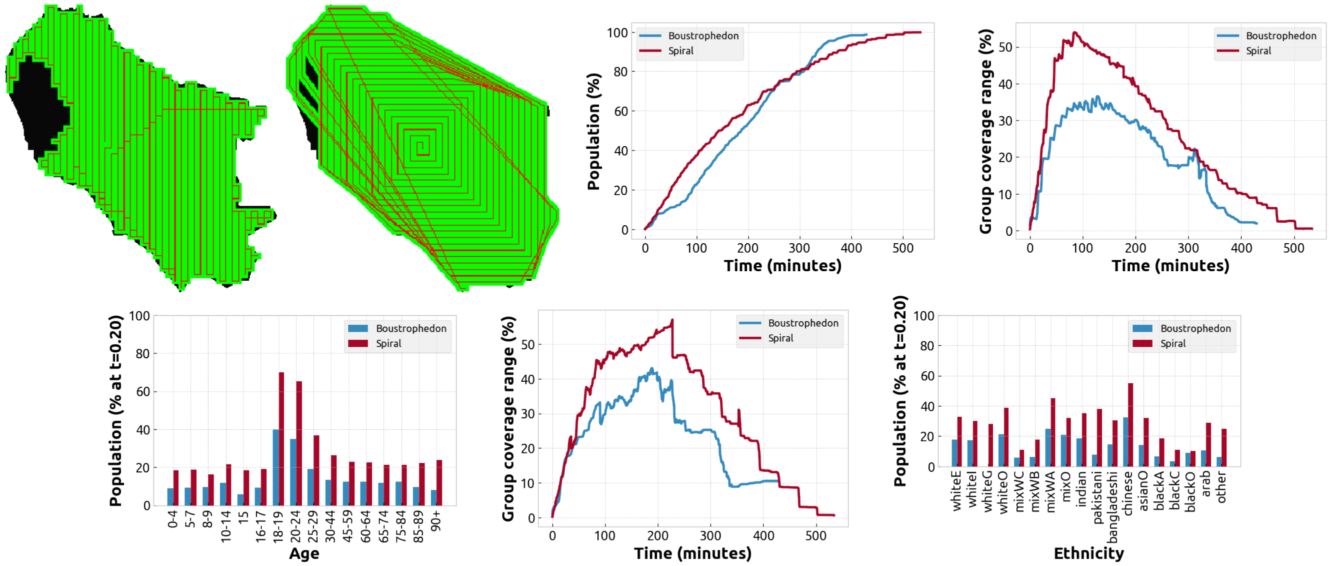

Fig. 1. Disparate coverage of age and ethnicity groups in the simplified unlimited-battery coverage planning problem. From left to right, top to bottom: Boustrophedon coverage plan, spiral coverage plan, cumulative coverage graph, range of coverage values across age groups (i.e. max-minus-min of the coverage percentage of all groups), the coverage values of all age groups at a particular time (0.20T), range of coverage across ethnicity, coverage of ethnicity groups at 0.20T.

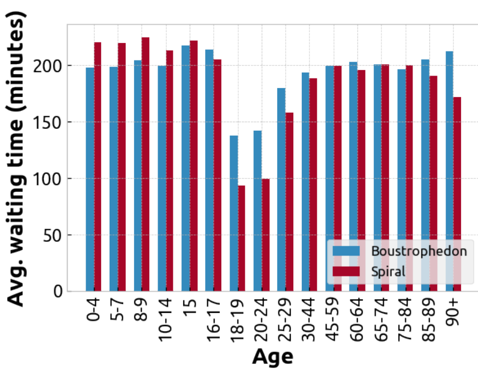

Fig. 2. Disparate waiting times across age groups in the simplified unlimited-battery coverage planning problem.

bias in quality of search and rescue missions to these very groups.

We also estimated average wait-times per age group, i.e. the average time it takes for a person of a certain group to be seen by the search-robot. Fig. 2 shows similar group biases to those seen in Fig. 1.

B. Limited-battery, low-waiting-time coverage

We now consider a more realistic search-and-rescue setting where the search-drone has limited battery and hence can only move for a maximum path length in each trip. We use 10km (500 cells) as the maximum path length as in other work [26], and assume each trip starts and ends in the same location (i.e. there is only one launching/charging station). Finally, the response team wishes to minimize average waiting time of the whole population, and hence the coverage algorithm explicitly prioritizes visits to highly populated areas. We consider two algorithms. The first is an adaptation of Boustrophedon: we split the map into N regions, and compute Boustrophedon coverage paths for each region. Then, the drone visits each region in descending order of region-population. The second algorithm is the evolutionary method of [9] with a large number of waypoints over multiple trips. To obtain low waiting times we set the maximization objective to the sum of the cumulative covered population (i.e. $\sum_{t=1}^{T} C_{m,\pi}(t)$ where $C_{m,\pi}(t)$ is the total population found along path π up to time t over map m; T is the length of path π). This objective promotes paths that have large covered populations from early on.

Fig. 3 shows the coverage paths, cumulative population coverage, and age-coverage fairness metrics. Paths have considerably more center-directed trips (i.e. radial red lines) since the drones need to return to the station for charging. Consistently with other work [7], the evolutionary method performs considerably better than classical methods—as can be seen by the population coverage graph. As in the simplistic unconstrained case, there are large degrees of inequality. Inequality peaks at around 20% of the coverage path as before and reaches 20 to 35% ranges (e.g. 60% coverage for young population vs 40% elderly). Although the inequality is slightly lower than in the unconstrained case, it could still represent considerable harm in terms of number-of-lives saved.

IV. Anytime-fair coverage method

We now propose an algorithm for anytime-fair coverage. We call "anytime-fair" to the goal of obtaining a distribution of group-coverage that is as fair as possible at any point in time. For the rest of the paper we will assume that a "fair" coverage is one where all social groups have been equally

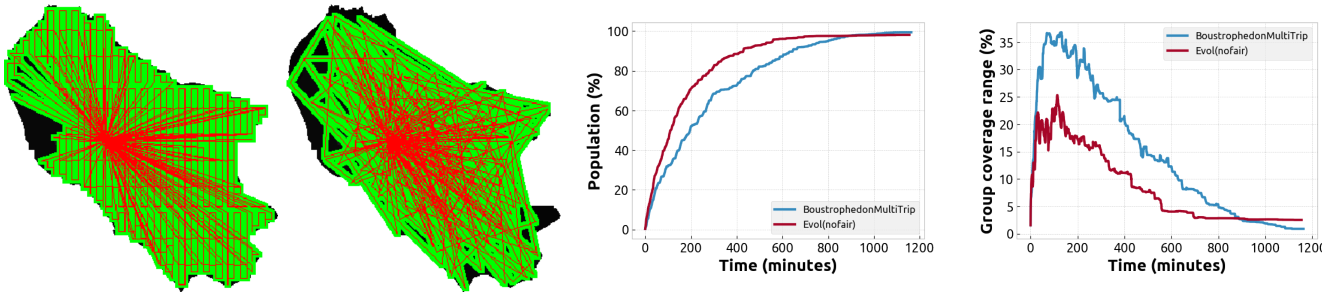

Fig. 3. Disparate coverage of age groups in the limited-battery coverage planning problem. From left to right, top to bottom: Boustrophedon coverage plan, evolutionary least-waiting-time coverage plan, cumulative coverage graph, range of coverage values across age groups (i.e. max-minus-min of the coverage percentage of all groups).

covered (e.g. coverage percentage of the populations in all age bins is the same). One issue with computing paths that are fair *exactly* is that this may be infeasible or come at a large efficiency cost, as recently shown for goal-directed path planning [9]. Therefore a trade-off needs to be found between the total population covered and coverage fairness.

We can formalize an anytime-fair method as one that minimizes the sum of cumulative unfairness: $\sum_{t=1}^{T} U_{m,\pi}(t)$, where $U_{m,\pi}(t)$ is an unfairness metric measured on map m along path $\pi(\tau)$ over the interval $\tau = [1, t]$. While optimizing this objective is trivial in an evolutionary-optimization setting such as that taken in previous goal-directed fair-planning work [9], such algorithms are slow to converge in the coverage-planning case due to high dimensionality of the search space (see Section V). In this paper we instead take a greedy approach to the problem that simultaneously maximizes average coverage and minimizes the range of coverage values across social groups, at all instants of time.

Algorithm 1: AnytimeFairCoverage

Input: map m, maximum trip length L
Output: coverage path π
$\pi = []$;
while *Coverage(π) < 100%* **do**
  $p =$ ExtractLastTripFromPath(π);
  $B = L - \text{Length}(p)$;
  $\pi =$ ExtendPathBestQuantile(m, π, B);
end

Basically, our method works by incrementally building the coverage path by sampling new straight-line traveling segments that maximize the mean minus standard deviation (i.e. 0.159-quantile assuming a normal distribution) of the group-coverage distribution. Algorithms 1 and 2 show the pseudo-code for our method. At each iteration the method greedily picks the next-best traveling segment out of a set of samples, incrementally until full coverage is reached (Algorithm 1). The greedy selection strategy is the core component of the method (Algorithm 2). The algorithm samples up to MAX_SAMPLES points (lines 2-3). To each sampled point x, the algorithm computes the straight-line segment path p from the end of the current path $\pi(T)$

Algorithm 2: ExtendPathBestQuantile

Input: map m, path-so-far π, budget B, sensor radius R, station s
Output: extended path π_{new}
1 $T = \text{Length}(\pi)$; $Q = []$;
2 **for** $i = 1, ..., $ *MAX_SAMPLES* **do**
3   $x =$ UniformSampleWithinBudget(m, B, $\pi(T)$, s);
4   **if** *not* x **then**
5     $p =$ BresenhamLineInterpolation($\pi(T)$, s);
6   **return** $\pi_{\text{new}} = [\pi, p]$
7   **end**
8   $p =$ BresenhamLineInterpolation($\pi(T)$, x);
9   $\pi' = [\pi, p]$;
10   $d =$ PercentagePeopleFoundInEachGroup(m, π');
11   $q =$ Mean(d) - Std(d);
12   $Q(\pi') = q$;
13 **end**
14 **return** $\pi_{\text{new}} = \text{argmax}_{\pi'} Q(\pi')$;

until x using Bresenham line interpolation, and appends it to the current path π (lines 8-9). The algorithm then computes the percentage of population that is found in each social group by the new path π', as well as a quantile q of that distribution (lines 10-11). The algorithm will select the segment that leads to a highest value for that quantile (lines 12-14), since this simultaneously maximizes average population coverage and minimizes the range of coverage values across groups. One important point of the algorithm is the sampling. Samples are taken uniformly within the points of the map m that still allow enough distance for the drone to return to the station. Therefore, the function UniformSampleWithinBudget(m, B, $\pi(T)$, s) samples points uniformly from m, and discards any points x for which the inequality $||x - \pi(T)|| + ||s - x|| < B$ does not hold (i.e. the distance to x plus the distance back to the station has to be within the current budget). When such a point is not found, the drone concludes the trip and returns to the station (lines 4-7).

In the next section we will see how this algorithm performs in practice compared to evolutionary methods and sampling-based baselines. We skip comparisons with the Boustrophedon method since, as shown in Section III, it performs worse than the evolutionary method in all respects.

V. Evaluation

We used the same setup as described in Section III to evaluate our method. As before we compute coverage paths for the city of Oxford, UK, from an hypothetical launch/charge station in the central fire station of the city. As baselines we use:

- "Evol": the evolutionary algorithm proposed in [9], adapted in the following way: 1) waypoints are scattered across multiple trips with forced station-returns at the end of each trip, 2) we compute the Pareto-front of two objectives: total coverage and unfairness ($\sum_{t=1}^{T} C_{m,\pi}(t)$ and $-\sum_{t=1}^{T} U_{m,\pi}(t)$). Therefore, this baseline optimizes anytime-fairness directly. The optimization is run until a plateau is reached (roughly 200 iterations);
- "Ours(nofair)": an adapted version of our method where segments are chosen in order to maximize total coverage $C_{m,\pi}(t)$ instead of the mean-minus-standard-deviation of group-coverage—thus ignoring the fairness component of the problem.
- "Ours(maxmin)": an adapted version of our method where segments are chosen in order to maximize the coverage obtained by the least-covered group (i.e. replacing line 11 in Algorithm 2 by $q = \min d$). This is a typical approach to fairness in planning [27], [14] and networking problems, sometimes called "Rawlsian" fairness [9] due to its similarity to John Rawls' theory of justice [28].

Both our method and its variants use 100 maximum samples (MAX_SAMPLES) in all experiments.

Figures 4 and 5 show the resulting coverage and unfairness graphs for the two cases of age and ethnicity fairness. Since the objective of the coverage algorithm is different in each case (minimize age- or ethnicity-unfairness), paths obtained by the algorithms are different in each case. Comparing the efficiency/speed of coverage ($\sum_{t=1}^{T} C_{m,\pi}(t)$) amounts to comparing the area under the curve of the coverage graph, while "anytime-fairness" ($\sum_{t=1}^{T} U_{m,\pi}(t)$) is the area under the curve of the "Group coverage range" graph.

The figure shows that our method achieves slightly higher coverage efficiency and higher final coverage than the evolutionary method: i.e. the evolutionary method fails to achieve total coverage. This is because the evolutionary method relies on chance (in the mating and mutation processes) to eventually cover the whole map with randomly-generated waypoints—making it hard to achieve full coverage. Our method also achieves a near-identical unfairness curve to "Evol", but performs slightly better since it reaches exactly 0 unfairness at the end of the path due to achieving full coverage. The fairness-unaware baseline (Ours-nofair) achieves similar coverage performance to ours but very high group bias: the range of group-coverage values peaks at 30% vs 10% for our method in the case of age-fairness (and 35% vs 20% in the case of ethnicity). These gaps are reflected in the per-group waiting time graph. The fairness-unaware method starts by covering highly-dense areas (of mostly 18-24 years old), thus decreasing the average waiting time for these groups. It does this at the cost of slower coverage for the elderly population (85+), while our method achieves a more balanced coverage (lower for the student population but higher for the elderly). The same happens in terms of ethnicity: a fairness-blind approach penalizes minorities such as Black and Arab, and privileges the dominant White ethnicity—while our approach is more balanced. As discussed in Section III this fairness property could be crucial for both increasing rescue effectiveness for the more at-risk population, as well as avoiding previous criticism of minority marginalization. Finally, and perhaps more importantly, our method obtains more balanced coverage of social-groups without significant changes in terms of total coverage efficiency. Particularly in the case of ethnicity-fairness the coverage curve graphs are coincident. This means that there is enough flexibility in the space of coverage plans to allow for fair-but-efficient coverage without heavy trade-offs. Using the popular maxmin strategy within our method ("Ours(maxmin)") leads to similar coverage and unfairness curves, although at slightly lower performance: this is most clearly visible in the average waiting time graphs, which show 154 minutes for maxmin vs 149 minutes for our method in the age-fairness case (Fig. 4) and 148 vs 144 minutes in the ethnicity-fairness case (Fig. 5).

We also estimated average wait-times per age group. Fig. 6 shows similar results to those seen in Fig. 4 and 5: our method achieves more balanced waiting times that are lower than "Ours(nofair)" for the older population and higher for the younger population—while being on average lower than "Evol". Using a maxmin strategy within our method ("Ours(maxmin)") leads to higher waiting times than our method for all social groups, as seen in the figure.

Besides achieving exact full coverage, our method has one more advantage compared to the evolutionary method of [9]: computation time. Table I shows the computation times for both methods in the age-group case. Times are averaged over 10 runs for our method, but over only two runs for the Evol method due to its extremely large values (it would take roughly 3 days to do 10 runs). The table shows that our method is more than 200 times faster than the evolutionary method, taking an average of 140 seconds to solve (vs 32375s). It also shows that the computational burden of promoting anytime-fairness is low, since it takes only 56s more than the fairness-unaware baseline—which is small within the timescale of total mission time (16h) and average waiting time (150min).

VI. Conclusion

In this paper we investigated an important social dimension of coverage path planning—fairness. We first demonstrated through realistic simulations of robot deployments that there is a fairness dimension to coverage planning. Fairness in this context is not related to who is seen or served

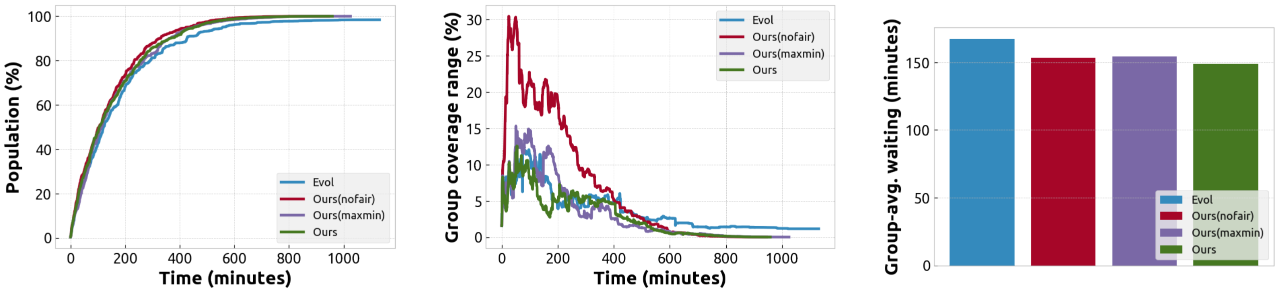

Fig. 4. Coverage, unfairness, and average waiting-time of our algorithm vs baselines. Age-group fairness optimization case. Waiting times are averaged over all age groups.

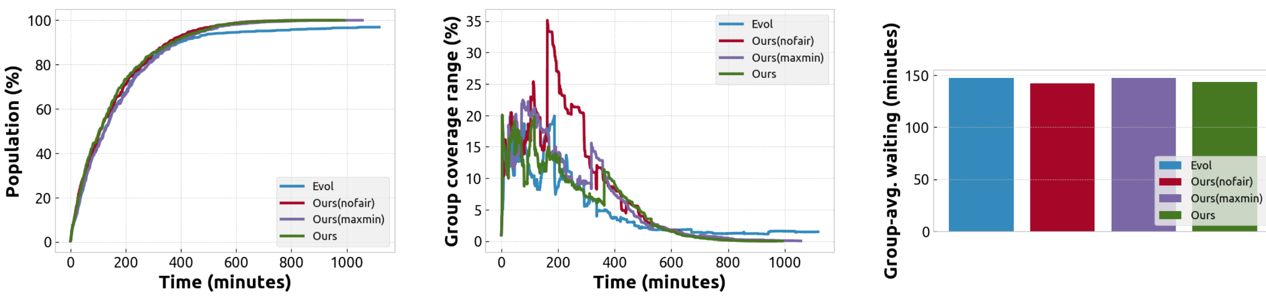

Fig. 5. Coverage, unfairness, and average waiting-time of our algorithm vs baselines. Ethnicity-group fairness optimization case. Waiting times are averaged over all ethnicity groups.

TABLE I
COVERAGE-PATH COMPUTATION TIME

Method	Planning time (s)
Evol	32375 ± 281
Ours (nofair)	84 ± 9
Ours	140 ± 15

by a robot—since the whole physical space of interest will eventually be covered. Instead, concerns of fairness could come from the *order* or *speed* with which different social groups are covered. We show that classical coverage algorithms, especially those that try to minimize average waiting times, will have biases related to the spatial segregation of social groups. These biases could be of concern since they may not be aligned with the interests of the application, e.g. disaster response.

We explored the example of a search drone in the context of disaster response. In this context, responses are expected to not privilege any social group, or at least privilege only those that are expected to be most at-risk. We showed that in the example of the city of Oxford, UK, a traditional search drone would privilege people in their early-20s, White, and Chinese. Such bias would be against equality and risk-prioritization-related principles of disaster response [2], [3], [4]. They would also lead to reinforcing existing inequalities and thus criticisms in disaster response, related to a lack of help to those social groups that are already marginalized [5].

To better align coverage planning with its fairness dimension we proposed a fairness-aware planning algorithm that is anytime-free. The concept refers to the objective of optimizing fairness at all instances of time. We show that our greedy sampling-based algorithm is 200 times faster to compute than existing evolutionary algorithms—while obtaining similarly fair and overall-faster coverage. We show the method's choice of optimizing the mean minus standard deviation of per-group coverages is also more effective than using a more traditional maxmin strategy within the method. Importantly, our algorithm successfully exploits the redundancy in coverage planning to achieve fairness at a low cost to efficiency.

Interesting directions of further research include extensions to multi-agent planning and station optimization. Additionally, we would like to investigate the performance of the algorithm in extremely segregated cities, as well as the potential of using the planning method to *characterize* segregation. Finally, we believe there is a need to close the responsible innovation loop by obtaining feedback from stakeholders in disaster response, for example through realistic user studies and workshops, and using this feedback to iterate the design of coverage planners.

REFERENCES

[1] B. M. Beamon and B. Balcik, "Performance measurement in humanitarian relief chains," *International Journal of Public Sector Management*, 2008.

[2] R. Hassanzadeh and Z. Nedovic-Budic, "Where to go first: prioritization of damaged areas for allocation of urban search and rescue (usar) operations (pi-usar model)," *Geomatics, Natural Hazards and Risk*, vol. 7, no. 4, pp. 1337–1366, 2016.

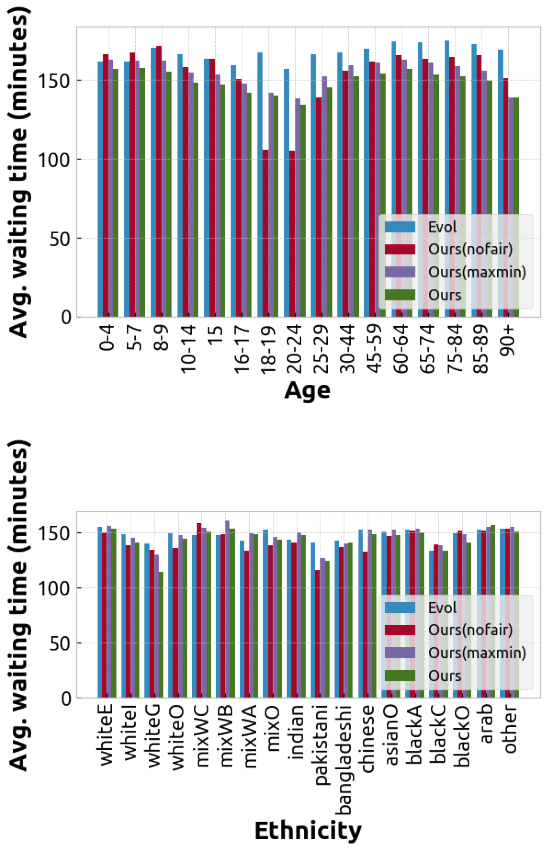

Fig. 6. Per-social-group waiting times obtained by our method vs baselines, in the age-fairness case (top) and ethnicity fairness case (bottom).

[3] G. Bognar and I. Hirose, *The Ethics of Health Care Rationing: An Introduction*. Routledge, 2014.

[4] O. Merin, N. Ash, G. Levy, M. J. Schwaber, and Y. Kreiss, "The israeli field hospital in haiti-ethical dilemmas in early disaster response," *New England Journal of Medicine*, vol. 362, no. 11, p. e38, 2010.

[5] B. Bolin and L. C. Kurtz, "Race, class, ethnicity, and disaster vulnerability," *Handbook of disaster research*, pp. 181–203, 2018.

[6] H. X. Pham, H. M. La, D. Feil-Seifer, and M. Deans, "A distributed control framework for a team of unmanned aerial vehicles for dynamic wildfire tracking," in *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2017, pp. 6648–6653.

[7] Y. Wang, T. Kirubarajan, R. Tharmarasa, R. Jassemi-Zargani, and N. Kashyap, "Multiperiod coverage path planning and scheduling for airborne surveillance," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 54, no. 5, pp. 2257–2273, 2018.

[8] I. Maza, F. Caballero, J. Capitán, J. R. Martínez-de Dios, and A. Ollero, "Experimental results in multi-uav coordination for disaster management and civil security applications," *Journal of intelligent & robotic systems*, vol. 61, no. 1, pp. 563–585, 2011.

[9] 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, 2020.

[10] H. Choset and P. Pignon, "Coverage path planning: The boustrophedon cellular decomposition," in *Field and service robotics*. Springer, 1998, pp. 203–209.

[11] 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.

[12] G. De Giacomo, F. Patrizi, and S. Sardina, "Generalized planning with loops under strong fairness constraints," in *Twelfth International Conference on the Principles of Knowledge Representation and Reasoning*, 2010.

[13] G. E. Fainekos, H. Kress-Gazit, and G. J. Pappas, "Hybrid controllers for path planning: A temporal logic approach," in *44th IEEE Conference on Decision and Control*. IEEE, 2005, pp. 4885–4890.

[14] P. Matl, R. F. Hartl, and T. Vidal, "Workload equity in vehicle routing problems: A survey and analysis," *Transportation Science*, vol. 52, no. 2, pp. 239–260, 2018.

[15] S. Jütte, D. Müller, and U. W. Thonemann, "Optimizing railway crew schedules with fairness preferences," *Journal of Scheduling*, vol. 20, no. 1, pp. 43–55, 2017.

[16] R. Baum, D. Bertsimas, and N. Kallus, "Scheduling, revenue management, and fairness in an academic-hospital radiology division," *Academic radiology*, vol. 21, no. 10, pp. 1322–1330, 2014.

[17] A. Golub, R. A. Marcantonio, and T. W. Sanchez, "Race, space, and struggles for mobility: Transportation impacts on african americans in oakland and the east bay," *Urban Geography*, vol. 34, no. 5, pp. 699–728, 2013.

[18] G. Walker, *Environmental justice: concepts, evidence and politics*. Routledge, 2012.

[19] M. Brandao, "Age and gender bias in pedestrian detection algorithms," in *Workshop on Fairness Accountability Transparency and Ethics in Computer Vision, CVPR*, Jun 2019.

[20] J. Buolamwini and T. Gebru, "Gender shades: Intersectional accuracy disparities in commercial gender classification," in *Proceedings of the 1st Conference on Fairness, Accountability and Transparency*, ser. Proceedings of Machine Learning Research, S. A. Friedler and C. Wilson, Eds., vol. 81. New York, NY, USA: PMLR, 23–24 Feb 2018, pp. 77–91.

[21] A. Chouldechova, "Fair prediction with disparate impact: A study of bias in recidivism prediction instruments," *Big Data*, vol. 5, no. 2, pp. 153–163, 2017.

[22] S. Barocas and A. D. Selbst, "Big data's disparate impact," *California Law Review*, vol. 104, p. 671, 2016.

[23] A. Singleton. (2015) 2011 census open atlas. university of liverpool: Consumer data research centre (cdrc). [Online]. Available: https://data.cdrc.ac.uk/product/cdrc-2011-census-open-atlas

[24] J. V. Hurtado, L. Londoo, and A. Valada, "From learning to relearning: A framework for diminishing bias in social robot navigation," *Frontiers in Robotics and AI*, vol. 8, p. 69, 2021.

[25] A. C. Kapoutsis, S. A. Chatzichristofis, and E. B. Kosmatopoulos, "Darp: divide areas algorithm for optimal multi-robot coverage path planning," *Journal of Intelligent & Robotic Systems*, vol. 86, no. 3-4, pp. 663–680, 2017.

[26] S. Choudhury, K. Solovey, M. J. Kochenderfer, and M. Pavone, "Efficient large-scale multi-drone delivery using transit networks," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2020, pp. 4543–4550.

[27] M. Malencia, V. Kumar, G. Pappas, and A. Prorok, "Fair robust assignment using redundancy," *arXiv preprint arXiv:2102.06265*, 2021.

[28] J. Rawls, *A Theory of Justice*. Harvard University Press, 1971.