

Socially-Driven Collective Path Planning for Robot Missions

Juan Camilo Gamboa Higuera, Anqi Xu, Florian Shkurti, and Gregory Dudek

School of Computer Science

McGill University, Montreal, Canada

{gamboa, anqixu, florian, dudek}@cim.mcgill.ca

Abstract—We address the problem of path planning for robot missions based on waypoints suggested by multiple human users. These users may be operating under distinct mission objectives and hence suggest different locations for the robot to visit. We formulate this problem using a constrained optimization approach by imposing various operational considerations, such as the robot’s maximum traversable distance. We then propose an approximative path planning algorithm with parameterized control over the degree of “social fairness” in the selection of waypoints from different users. Through a user study, we compared the performance of the proposed path planner under different fairness settings and for different mission scenarios.

Keywords—waypoint-based path planning; multi-human-single-robot system; constrained optimization; social fairness.

I. INTRODUCTION

We propose a robot path planning algorithm that achieves coverage of waypoints within a bounded region, where waypoints are issued by multiple human users. The robot’s path is constructed using one of several waypoint selection strategies, while adhering to various operational constraints such as the maximum traversable distance within a single session. Our solution aims to satisfy all users, even for cases where the provided sets of waypoints may differ drastically from one another. This path planner is suitable for situations where the robot cannot visit all of the specified locations in a single session; this may arise when the combined coverage region well exceeds the robot’s operational range, or when the users are driven by conflicting mission objectives.

This type of supervisory tele-robotic setup can be potentially applied to a wide variety of applications, most notably including search and rescue, and planetary exploration. In these contexts, it is often practical to delegate the task of identifying mission objectives to human experts. For instance, our research group is interested in applying this path planner to a heterogeneous multi-robot system for assisting marine biologists in the study of coral reefs. Specifically, our robotic aircraft carries out coverage of the target reef region [1] while streaming live aerial footage to off-site scientists. These scientists then identify locations for further inspection, based on their expert domain knowledge and field experience of coral reefs. Using the proposed planning algorithm, we can compute a mission path for our underwater robot [2] to collect up-close footage at as many

of these sites as possible during a single session, so as to deliver the data back to the biologists in a timely manner.

In this work we present mathematical formulations for the length-constrained path planning problem using waypoint input from multiple users. We then show that these *Constrained Coverage Planning* and *Discrete Constrained Coverage Planning* problems are NP-hard. This motivates the primary contribution of our work, which is a socially-driven approximative planning algorithm for covering as many of the users’ points as possible while adhering to mission-critical constraints, namely the maximum length of the trajectory, and the robot’s given starting and ending locations. A key novelty of this planner is its parametric control over the degree of *social fairness* of the generated path, in other words, whether the robot’s trajectory will cover an evenly distributed or imbalanced subset of waypoints from each user. More specifically, our planner can be configured to select waypoints based on the following strategies:

- *elitism*: where the path covers locations submitted by a subset of users predominantly;
- *egalitarianism*: where the path visits an almost equal number of locations selected by each and every user;
- *nearest neighbor*: where the goal is to simply cover as many nearby waypoints as possible.

We further present an evaluation of this collective path planning algorithm through a user study. In this study, participants first select waypoints over a bounded terrain map, while adhering to one of several mission scenarios. In a follow-up stage, users provide relative preferential rankings for robot paths generated using multiple parameterizations of our socially-driven planning algorithm. These rankings are then used to compare the performance of our waypoint selection strategies for a variety of mission setups.

II. RELATED WORK

Many research groups have studied human-robot systems where the task of identifying the robot’s mission objectives is delegated, or *crowdsourced*, to human experts. For instance, Bitton and Goldberg [3] proposed a mixed-initiative control system for multiple Unmanned Aerial Vehicles (UAV) to conduct search of a stationary target. Their simulation framework allows human operators and automated agents to collaboratively define search locations within a bounded

region. Their algorithm then computes an optimal target-to-UAV assignment by minimizing the total number of search iterations. Another related work involves a multiple-operator-single-robot setup where the robot is controlled via a spatial voting scheme based on live camera feedback [4]. The authors applied an interval clustering method to fuse waypoints from different humans into a single aggregated robot command. One possible outcome of this technique is that the robot may not ever satisfy requests from a sub-group of users, which is an aspect that is addressed by our current work.

Cummings *et al.* [5] studied the benefits of allowing a human operator to adjust robot paths that are generated using a planning algorithm, within the context of lunar exploration. Their user study revealed that the improvements to the path quality varied depending on the evaluation criterion used by the human, be it the duration, distance, terrain slope, or fuel limitations of the path.

A similar body of work addresses the problem of shared control over a pan-tilt streaming video camera, which has been studied by Song *et al.* [6] and by Liu *et al.* [7]. Furthermore, several groups [8], [9] have evaluated crowd-sourcing as a viable alternative to solving challenging visual recognition problems within robotic systems.

Finally, the collective path planning problem in this work is related to multi-armed bandit problems [10], [11] and exploration-exploitation-style planning strategies that attempt to maximize the utility of the bandit or path given a limited number of trials. A key difference however is that these problems in general allow the plan to be revised based on updated observations, whereas our formulation leads to a fixed path that is carried out by the robot without alterations. Similar off-line planning problems have been studied in the area of constrained network routing [12], and, in addition, various randomized planning algorithms [13] have been applied to off-line exploration-exploitation schemes [14].

III. COLLECTIVE PATH PLANNING PROBLEMS

We consider a multiple-user-single-robot setup where the robot's role is to visit and collect data at various locations of interest using an omni-directional sensor. Given a map of the bounded operational region, each human user is tasked with identifying locations of interest for the robot to visit, based on a given mission objective that might differ among users. As a result, these sets of locations could potentially be dissimilar and far from each other, due to varying selection criteria, priorities, and preferences. Despite this, we would like to compute a path that allows the robot to cover as many of these locations as possible within a single session.

A. Constrained Coverage Planning

We begin by presenting a generalized version of the problem, which we refer to as *Constrained Coverage Planning* (CCP). Formally, we assume that users $i = 1..m$ have

corresponding interest functions $I_i(x, y)$, whose range is $[0, 1]$ over a common and bounded 2-D Cartesian map M . In order to compare different users' interests, we also require that these functions be normalized:

$$\iint_M I_i(x, y) dx dy = 1 \quad \text{for all } i = 1..m$$

Next we denote s and t as the starting and ending positions of the robot's path, which are assumed to be given as part of the mission setup. We also assume that the robot is equipped with an omni-directional sensor with fixed radius r .

The objective of the CCP problem is to find a path P in Cartesian space that satisfies the following conditions:

- (I) the collective user interest covered by P is maximal among all possible paths from s to t ;
- (II) P starts at s and ends at t ; and
- (III) the Euclidean arc length of the path P is less than a given maximum distance threshold L .

To express (I) mathematically, we define P_r as the area around the path obtained by dilating¹ P by r . Finding a path P that maximizes the collective user interest thus corresponds to solving the following optimization problem:

$$\operatorname{argmax}_P \sum_{i=1}^m \iint_{P_r} I_i(x, y) dx dy \quad (1)$$

In order to generate a feasible path that the robot can execute within a single session, our optimization objective must adhere to pragmatic constraints, namely (II) and (III). The former condition covers setups where the robot is exploring a region while on route to a final destination, as well as other missions where the robot surveys its surroundings and then returns to the home base. Constraint (III) is motivated by the limited battery life of certain types of robotic vehicles. We make the simplifying assumptions that the omni-directional sensor is operating continuously and that the robot is moving at a fixed speed. Consequently, we can empirically derive the maximum distance, L , traversable by a given robotic vehicle.

B. Discrete Constrained Coverage Planning

In order to solve Eq. 1, each user must provide their interest values $I_i(x, y)$ over all positions within the bounded operational region. In general this process is tedious and impractical, especially for cases where there are many concentrated sub-regions with high levels of interest. We thus propose the *Discrete Constrained Coverage Planning* (D-CCP) variant of the problem, where instead of querying for $I_i(x, y)$ directly, we elicit from each user $i = 1..m$ an unordered set of points S_i that corresponds to a sample-based approximation of their underlying interest function $I_i(x, y)$. This sample-based approximation can be expressed as:

¹In other words, P_r is the Minkowski sum of P with a disk of radius r .

$$I_i(x, y) \approx \begin{cases} \frac{1}{|S_i|} & \text{if } (x, y) \in S_i \\ 0 & \text{otherwise} \end{cases} \quad (2)$$

This discretization allows us to formulate D-CCP as a graph-based planning problem. Specifically, we consider a complete graph $G = (V, E)$ overlaid on top of the Cartesian map, where the vertices are $V = \{s, t\} \cup S$ for $S = \cup_i S_i$, and where edges are weighted by the Euclidean distance between their two endpoints. For instances of the problem where the robot's sensing range happens to be significantly smaller than the distance between vertices in G , the optimization formulation in Eq. 1 then reduces to finding a path P on G that visits the maximum number of unique nodes in V , while adhering the constraints (II) and (III).

We now define $\mathbf{1}_{[v \in P_r]} \in \{0, 1\}$ as the indicator variable that is equal to 1 iff node v is covered by P_r , and we also define $\mathbf{1}_{[(u, v) \in P]} \in \{0, 1\}$ to be 1 iff edge $(u, v) \in P$. Furthermore, we denote $d(u, v)$ to be the Euclidean distance between the points $u, v \in V$. These are used in the constrained optimization formulation of D-CCP, as shown below:

$$\begin{aligned} \max_P \quad & \sum_{i=1}^m \sum_{v \in S_i} \mathbf{1}_{[v \in P_r]} \\ \text{subject to} \quad & \sum_{(u, v) \in E} \mathbf{1}_{[(u, v) \in P]} d(u, v) \leq L; \text{ and} \\ & P \text{ is a simple path from } s \text{ to } t. \end{aligned}$$

By considering only instances where points not in the piecewise-linear path P are not covered by its dilated counterpart P_r , the exact solution can be potentially computed using an integer program. One of the key challenges with this approach is to express the connectivity of the path using a polynomial number of constraints, so that the construction of the problem itself constitutes a feasible computation. Though this issue has been addressed [15], [16] previously, we chose not to follow the integer programming paradigm since it does not guarantee polynomial running time.

C. Complexity Analysis

The following theorem is the main motivation for the approximative planning algorithm that we shall present in Section IV.

Theorem 1. *Both the generalized and discrete Constrained Coverage Planning problems are NP-hard.*

Proof: Our proof involves a Turing reduction from a version of the Traveling Salesman Problem to D-CCP and then to CCP. It has been shown that the decision version of the discretized Euclidean Traveling Salesman Problem (DE-TSP) is NP-complete [17]. DE-TSP addresses the question of whether there exists a Hamiltonian cycle of an *integer* length of at most L that moves through a set of *integer-valued* points, S_i , based on a modified Euclidean distance metric $d(x_1, x_2) = \lceil \|x_1 - x_2\|_2 \rceil$.

Any arbitrary instance of DE-TSP can be mapped into an instance of D-CCP or CCP in polynomial time. This is achieved by first considering one of the points as the robot's starting and ending locations, i.e. $v = s = t$ for some $v \in S$. Next, we set the sensor radius r to be sufficiently small so as to prevent implicit coverage of the additional points as the robot moves from one point to the next in the path, i.e.:

$$r < \operatorname{argmin}_{a, b, c \in S} \left\{ \operatorname{Dist}(a, \overrightarrow{b-c}) \right\}$$

where $\operatorname{Dist}(a, \overrightarrow{b-c})$ measures the point-to-line-segment Euclidean distance, and where the $\operatorname{argmin}\{\cdot\}$ excludes co-linear point triplets. This can be computed in $O(|S|^3)$ time.

We can further reduce the DE-TSP instance into the general CCP formulation by defining $I_i(x, y)$ according to Eq. 2. DE-TSP can thus be solved by verifying if the exact solution to D-CCP or CCP for a given L covers all available points. Since we have shown that CCP and D-CCP are Turing-reducible to an NP-complete problem in polynomial time, we conclude that they are both NP-hard. ■

IV. SOCIALLY-DRIVEN COLLECTIVE PATH PLANNER

We now present a greedy algorithm that produces an approximative solution to the Discrete Constrained Coverage Planning (D-CCP) problem. This iterative algorithm begins by setting a *candidate path* as the straight line from the robot's starting and ending locations. Then, during each cycle of the path construction process, waypoints that are not in the candidate path are assigned utility scores, to be explained below. Our algorithm repeatedly selects the point with the highest utility and merges it into the candidate path, until the maximum distance threshold L is reached.

These utility scores reflect a parameterized degree of *social fairness* by taking into account the following factors:

- In general, to ensure coverage of the most number of points, the increase in path length during each iteration should be minimal. Therefore, points that are farther away from the current candidate path should have lower utility scores than those that are closer.
- If we prefer the equal inclusion of points from different users, then higher utility scores should be assigned to points belonging to users that have few points already included in the current candidate path.
- If we instead prefer to incorporate as many points as possible from a subset of users into the path, then the algorithm should assign higher utility scores to their points that are closer to the current candidate path.
- If we do not consider the fair or biased distribution of each user's points, then the utility scores should be inversely proportional to the distance between each point to the candidate path.

Based on these guidelines, if in the k -th iteration the algorithm has a candidate path P_k , then for the next iteration we

define the following utility assignments for the remaining unvisited points $v \in S_i$ belonging for a given user i :

$$U(v) = \left(\frac{|S_i \cap P_k|}{\sum_{i=1}^m |S_i|} \right)^{-f} \frac{1}{\text{Dist}(v, P_k)} \quad (3)$$

where $\text{Dist}(v, P_k)$ is the Euclidean distance² between the point v and the piecewise-linear candidate path P_k .

This utility scoring scheme is configurable via a *social fairness* parameter f : when $f = 0$, the algorithm behaves in a “nearest neighbor” (NN) setting, where the next point to be inserted is the one closest to the current candidate path, regardless of which user it comes from. When $f > 0$, then the choice of the next point will be among those from users with the least number of points already included in the candidate path, provided that the distances to said points are small. This parameter setting reflects an “egalitarian” (EG) strategy. In contrast, when $f < 0$ then users with the most points already included in the candidate path will be preferred, provided that their unvisited points are close to the candidate path; this reflects an “elitist” (EL) strategy.




Figure 1. Ratio of the utility scores $U(v_2)/U(v_1)$ for two users, plotted against the social fairness parameter f and the ratio of the closest-point-to-path distances d_1, d_2 . The solid purple curve depicts the decision boundary for which user’s point is chosen by our algorithm.

Fig. 1 illustrates which of two users’ closest waypoint to the candidate path is selected by our algorithm as we vary the fairness parameter f and the ratio of the distances d_1, d_2 of these waypoints. In this setup, we assume that the first user has more waypoints already in the candidate path than the second user. The decision boundary is located at $d_1/d_2 = 1$ for the nearest neighbor setting ($f = 0$), and this boundary is shifted to be more or less favorable to the user with the fewer number of chosen waypoints as f increases or decreases from 0.

²Also, to handle cases where $|S_i \cap P_k| = 0$, in practice we add a small ϵ to both the numerator and denominator in Eq. 3 to avoid division by zero.

Algorithm 1 Socially-driven collective path planner

Inputs: Disjoint sets $S_1, S_2, \dots, S_m \in S$ of each user’s points; robot’s starting position s and ending position t ; maximum path length L , where $L \geq \|s - t\|_2$; and degree of fairness $f \in \mathbb{R}$.

Output: An ordered sequence of points P from s to t .

```

 $P \leftarrow (s, t)$ 
pathAugmented  $\leftarrow$  true
 $V = S \cup \{s, t\}$ 
while pathAugmented do
  pathAugmented  $\leftarrow$  false
   $\Phi \leftarrow V \setminus P$ 
  while  $\Phi \neq \emptyset$  do
     $\phi \leftarrow \underset{v \in \Phi}{\text{argmax}} U(v)$ 
    if  $\text{length}(P \cup \{\phi\}) \leq L$  then
       $P \leftarrow P \cup \{\phi\}$ 
      pathAugmented  $\leftarrow$  true
    break
   $\Phi \leftarrow \Phi \setminus \{\phi\}$ 
return  $P$ 

```


V. USER STUDY

We conducted a user study to determine which way-point selection strategy would result in the highest user satisfaction, in terms of maximizing coverage based on their assigned mission objectives. We hypothesize that no single strategy will always out-perform others in all possible scenarios, since we suspect that the performance of a given planner configuration (i.e. value for f) will depend on the similarities of the different users’ interest functions $I_i(x, y)$.


Various characteristics of a mission setup can affect the similarity of different users’ $I_i(x, y)$. One key characteristic is the spatial concentration of the objective regions: if the most relevant regions for a given objective are concentrated within small and isolated areas, then the users’ waypoints would be likely to be spatially similar. On the other hand, if these regions span across relatively large areas, then it is less certain whether users would select waypoints that are close to each other. Another prominent factor is the specificity in the wording of the task objective. In particular, a precise and objectively-defined objective criterion would be more likely to lead to sets of similar waypoints, in contrast to a vaguely-defined objective that would require interpretation from each user. Furthermore, situations where the robot is time-shared by several groups of users with different agendas may also result in disparate waypoint sets.

A. Experiment Setup

To determine the effects of the aforementioned mission characteristics on the performance of our waypoint selection strategies, we devised the following missions:



(a) *Life on Mars* scenario



(b) *Coral Reef Survey* scenario

Figure 2. Waypoints submitted by 44 users, depicted as distinctly-colored circles, (a) for the *Life on Mars* scenario, where the objective is to inspect *large craters*; and (b) for the *Coral Reef Survey* scenario, where the goal is to identify *colorful corals*. The Mars satellite map is courtesy of Google Earth and partners (ESA/DLR/FU Berlin, Image NASA / USGS); the stitched reef photo is courtesy of Katrine Turgeon. (Note: color images.)

- (1) *Life on Mars*: Given a satellite map of a region on Mars, users are tasked to pinpoint locations for a rover to explore, in the search for signs of life. Users are told to assume that there is higher likelihood of finding life signs in *large craters*.
- (2) *Coral Reef Survey*: Given a stitched map of a coral reef, the mission objective is to catalog fish species near *colorful corals* using an autonomous underwater vehicle with a downward-facing camera.
- (3A) *Search and Rescue*: Users are informed of an injured person who is on the roof of a *tall building* within an urban area. The objective is thus to identify possible locations for an unmanned aerial vehicle to inspect, in search of this individual.
- (3B) *Insect Survey*: While provided with the same urban map as in (3A), users are asked to identify *green spaces* such as trees and terraces, so that a quadrotor can catalog the insect population near these locations.

We conducted a number of trials where we asked users to mark 10 locations of interest over a given map. We then divided the user base into two groups, and fed the waypoints for each user group into our socially-driven planner. Specifically, our algorithm generated paths using 4 fairness settings: *egalitarian* (EG, $f = 40$), *nearest neighbor* (NN, $f = 0$), *nearest elitist* (NE, $f = -2$), and *elitist* (EL, $f = -40$). Whereas the first two pair of trials involved missions (1) and (2) respectively, the third pair of trials was designed specifically to elicit distinct interest functions $I_i(x,y)$, by combining waypoints from half of the users for mission (3A) and the other half from (3B). We subsequently asked users to rank the four generated paths based on the objective of their mission for each of the three scenarios involved. This user study was carried out through a web-based interface for selecting waypoints and for ranking paths.

B. Results and Analysis

Our study involved 44 participants in total that were divided randomly into two groups of 22 users each. Each of the two groups was asked to evaluate paths for three different scenarios, and hence we obtained 132 path rankings in total.

As can be seen in Fig. 2, the chosen waypoints from the majority of participants are concentrated near distinct regions perceived to be most relevant to the mission objectives. In addition, the clusters of points in Fig. 2(a) are more localized than those in Fig. 2(b), which is consistent with our hypothesis that a concise and objectively-defined mission criterion (i.e. *large craters*) is more likely to result in similar interest functions $I_i(x,y)$ as opposed to a more interpretive criterion (i.e. *colorful corals*).

Fig. 3 shows the paths generated using the 4 fairness settings for the third scenario, which involved both missions (3A) and (3B). We observe that paths generated using the *nearest neighbor* and *nearest elitist* strategies are the least spread out from the robot's starting and ending locations, which were set to the center of the map in this user study. Also, the *egalitarian* (EG) path in Fig. 3(a) covers waypoints from a large number of users, though some of users had more waypoints selected than others based on their proximity to the evolving path. Furthermore, Fig. 3(d) illustrates that the *elitist* path primarily covered roofs based on the waypoints from a single user with the *Search and Rescue* scenario (3A), and thus spanned across the entire operational region. In contrast, the other generated paths visited locations of interest for both missions (3A) and (3B), and hence were less spread out.


Our results are in the form of lists of preferences, each of which reflects a strictly ordered ranking (i.e. one without ties) to the paths generated using different waypoint selection strategies.




(a) Path generated by an *egalitarian* (EG, $f = 40$) strategy



(b) Path generated by a *nearest neighbor* (NN, $f = 0$) strategy



(c) Path generated by a *nearest elitist* (NE, $f = -2$) strategy



(d) Path generated by an *elitist* (EL, $f = -40$) strategy

Figure 3. Paths generated by 4 fairness configurations of our socially-driven planner, for the mixed scenario of “Search and Rescue” (3A) and “Insect Survey” (3B). Each dilated path covered waypoints from multiple users, which are differentiated by the color of the markers. The starting and ending locations in this trial are both located at the center of the map and are not depicted. The 3-D satellite map is courtesy of Google Earth and DigitalGlobe.

To evaluate the overall ranking for the different strategies, we computed their average ranking positions and corresponding standard deviation values across users within each trial. In cases where all the standard deviations are small, then the ordering established by each strategy’s mean position provides a good indication of the aggregated user preference.

For cases where the users provided drastically different rankings, we determined the overall ordering using the *Kemeny-Young ranking* rule [18]. This rule generates an aggregated ranking by maximizing the number of pairwise agreements between the solution and each user’s ordering. The Kemeny-Young ranking also minimizes the sum of pairwise disagreements with all the preference lists, known as the *Kendall tau distance*, which is useful in reflecting the aggregate response across users. Furthermore, this ordering has the desirable property of satisfying *Condorcet’s criterion*: if a specific waypoint selection strategy is preferred over all others through pairwise comparisons, then it would be ranked in the highest position overall.

We also applied the Friedman test to each set of rankings for a given trial in order to assess the variability among the rankings from different users. When the resulting p value

is smaller than a reasonable threshold (e.g. $p < 0.05$) we then consider the aggregated ranking to reflect consistent behavior across users in a statistically significant manner.

Table I presents the average and Kemeny-Young aggregated rankings and the corresponding Friedman test p -values for each of our six 22-user trials. The Kemeny-Young rankings revealed no clear preference between the egalitarian (EG) and elitist (EL) strategies for trials 1-22-1, 1-22-2, 2-22-1 and 2-22-2. This is mainly attributed to the fact that users selected very similar locations for missions (1) and (2), as can be seen in Fig. 2(a) and Fig. 2(b). Since these results are all statistically significant (i.e. $p < 0.05$), we deduce that egalitarian (EG) and elitist (EL) strategies are likely to be preferred in situations where users share similar opinions on regions of interest. In these cases, our results also suggest that the nearest neighbor (NN) and near elitist (NE) strategies are consistently less preferred.

For the two trials involving the third scenario, which combined waypoints from missions (3A) and (3B) (i.e. 3-22-1 and 3-22-2), the elitist (EL) strategy emerged as the most preferred among our users. The egalitarian (EG) strategy, on the other hand, was not ranked consistently, though these results were not statistically significant (i.e. $p > 0.05$).

Table I
AVERAGE AND AGGREGATED PATH RANKINGS, ALONG WITH
FRIEDMAN'S TEST FOR STATISTICAL SIGNIFICANCE, FOR SIX 22-USER
TRIALS AND FOR THREE COMBINED 44-USER DATASETS.

[Scenario]- [# Users]-[Trial]		Relative Ranking				Friedman p-value
		(Best)			(Worst)	
1-22-1	Mean	EL	EG	NE	NN	< 0.01
	K-Y	EL	EG	NE	NN	
1-22-2	Mean	EL/EG	EL/EG	NE	NN	< 0.01
	K-Y	EG	NE	EL	NN	
2-22-1	Mean	EG	NN	EL	NE	< 0.05
	K-Y	EG	EL	NN	NE	
2-22-2	Mean	EL	EG	NE	NN	< 0.01
	K-Y	EL	EG	NE	NN	
3-22-1	Mean	NN	EL	EG	NE	0.06
	K-Y	EL	EG	NE	NN	
3-22-2	Mean	EL	NN	NE	EG	0.27
	K-Y	EL	NE	NN	EG	
1-44-1	Mean	EL	EG	NE	NN	< 0.01
	K-Y	EL	NN	NE	EG	
2-44-1	Mean	EL	EG	NN	NE	< 0.01
	K-Y	EL	EG	NE	NN	
3-44-1	Mean	NN	EL	NE	EG	0.06
	K-Y	EL	NN	NE	EG	

K-Y: Kemeny-Young aggregated ranking; NN: Nearest Neighbor;
NE: Nearest Elitist; EL: Elitist; EG: Egalitarian.

This meant that different users were not consistent in their preference orderings, which we anticipated given that half of the participants were interested in paths covering rooftops while the other half preferred paths near trees and grass terraces. Whereas the elitist strategy produced paths that addressed either mission (3A) or (3B) predominantly, the other strategies lead to paths that attempted to satisfy the two conflicting objectives simultaneously. Our results indicate that users in general did not prefer the latter set of paths.

The histogram plots in Fig. 4 depict the distribution of user preferences for each of the three scenarios, based on preference lists from all 44 participants. In conjunction with the corresponding orderings shown in Table I (i.e. 1-44-1, 2-44-1, and 3-44-1), we summarize that the elitist (EL) strategy was empirically shown to be the most preferred in almost all the setups. Based on this finding, we observe that an emphasis on pleasing a subset of users lead to the most overall user satisfaction. One possible reasoning is because the chosen waypoints are likely to be representative of a number of other users' interests, and thus participants who did not have their waypoints selected may still be relatively satisfied with the generated path. In addition, the egalitarian (EG) strategy exhibited similar performance compared to the elitist setting for scenarios (1) and (2), though the egalitarian-based paths were among the least favored in scenario (3), as seen from Fig. 4(c). We thus conclude that in a mixed-objective situation, fully pleasing some of the users completely is more rewarding than partially pleasing all of the users.

Our results also indicated that the nearest neighbor (NN) and nearest elitist (NE) settings produced paths that were almost always least preferred among the four planner con-

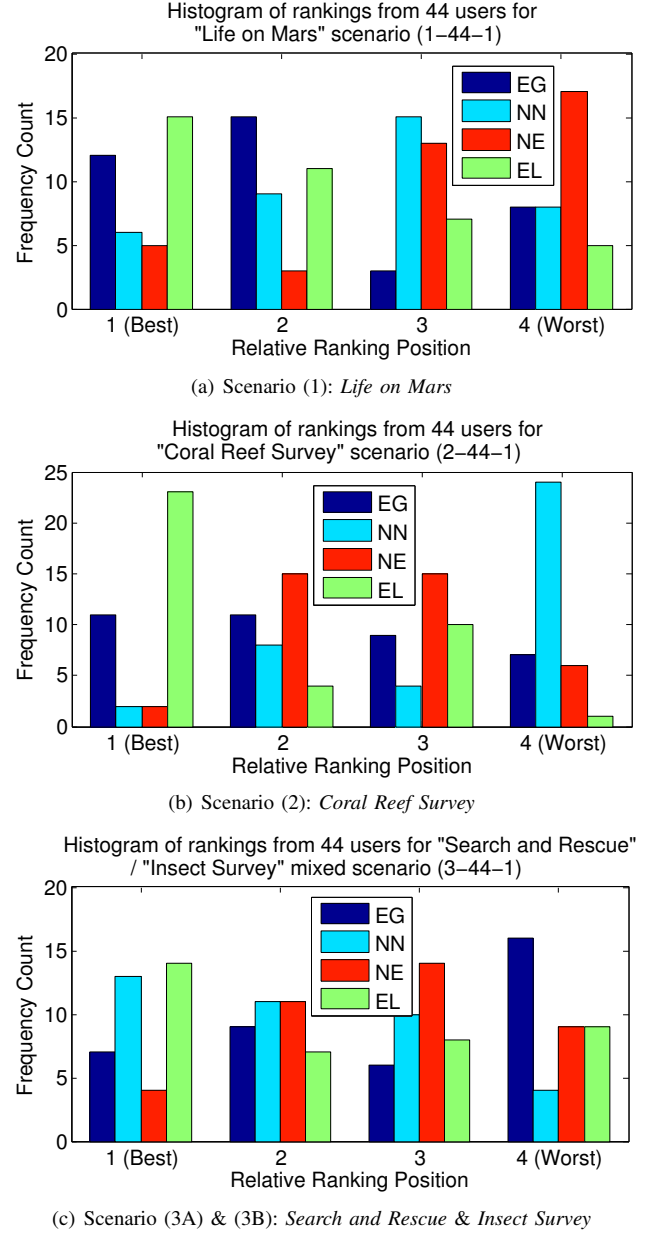


Figure 4. Histogram of rankings of each of the four algorithms by 44 users for the 3 different scenarios.

figurations. Although the nearest neighbor approach constructed paths that visited the most waypoints within the robot's operational range, user rankings indicated that these paths were not particularly effective at achieving coverage of the mission objectives. Our results were influenced in part by the use of a very restricted maximum traversable lengths L , which ensured that the generated paths could not perform dense coverage of the entire operational region. This caused the paths generated using the nearest neighbor strategy to be almost always concentrated around a small locality around the robot's starting and ending locations. We thus conclude

that for mission configurations where the robot's maximum traversable distance is limited relative to the size of the operational region, the nearest neighbor strategy will achieve less coverage across the entire region when compared to the elitist and egalitarian strategies.

VI. CONCLUSION

In this paper we investigated a multi-objective length-constrained path planning problem for human-assisted robot mission planning. Specifically, we proved that both *Constrained Coverage Planning* (CCP) and *Discrete Constrained Coverage Planning* (D-CCP) are NP-hard problems. Consequently, we proposed an approximative planning algorithm, which uses waypoints provided from multiple human users to compute a path that attempts to maximize the coverage of locations of interest towards one or more mission objectives, while adhering to the robot's maximum traversable distance. This iterative planner also incorporates an adjustable parameter that determines the degree of social fairness in its waypoint selection strategy. We further carried out a user study to examine the performance of these waypoint selection strategies in different mission scenarios. Our results indicate that an elitist selection strategy was most favorable among the majority of users our three scenarios, while, in contrast, the nearest neighbor strategy resulted in the least preferred paths overall.

In future work, we plan to investigate the self-configuration of the degree of fairness in our planner, based on a history of generated paths on multi-robot missions that are repeated in the same environment consecutively for many days. We are also motivated to integrate this planner into our heterogeneous multi-robot system for assisting marine biologists in monitoring coral reefs.

ACKNOWLEDGMENTS

We would like to thank Professor David Avis for contributing to our NP-hardness proof, and also to all the participants of our user study.

REFERENCES

- [1] A. Xu, C. Viriyasuthee, and I. Rekleitis, "Optimal complete terrain coverage using an unmanned aerial vehicle," in *IEEE Int. Conf. on Robotics and Automation (ICRA '11)*, 2011, pp. 2513–2519.
- [2] G. Dudek, M. Jenkin, C. Prahacs, A. Hogue, J. Sattar, P. Giguère, A. German, H. Liu, S. Saunderson, A. Ripsman, S. Simhon, L. A. Torres-Mendez, E. Milios, P. Zhang, and I. Rekleitis, "A visually guided swimming robot," in *IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS 2005)*, 2005, pp. 3604–3609.
- [3] E. Bitton and K. Goldberg, "Hydra: A framework and algorithms for mixed-initiative UAV-assisted search and rescue," in *IEEE Int. Conf. on Automation Science and Engineering (CASE '08)*, 2008, pp. 61–66.
- [4] K. Goldberg, D. Song, Y. Khor, D. Pescovitz, A. Levandowski, J. Himmelstein, J. Shih, A. Ho, E. Paulos, and J. Donath, "Collaborative online teleoperation with spatial dynamic voting and a human "tele-actor"," in *IEEE Int. Conf. on Robotics and Automation (ICRA '02)*, 2002, pp. 1179–1184.
- [5] M. Cummings, J. Marquez, and N. Roy, "Human-automated path planning optimization and decision support," *Int. J. Human-Computer Studies*, vol. 70, no. 2, pp. 116–128, 2012.
- [6] D. Song, N. Qin, and K. Goldberg, "Systems, control models, and codec for collaborative observation of remote environments with an autonomous networked robotic camera," *Autonomous Robots*, vol. 24, pp. 435–449, 2008.
- [7] Q. Liu, D. Kimber, J. Foote, L. Wilcox, and J. Boreczky, "FlySPEC: A multi-user video camera system with hybrid human and automatic control," in *10th ACM Int. Conf. on Multimedia*, 2002, pp. 484–492.
- [8] A. Sorokin, D. Berenson, S. Srinivasa, and M. Hebert, "People helping robots helping people: Crowdsourcing for grasping novel objects," in *IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS '10)*, 2010, pp. 2117–2122.
- [9] S. Chernova, N. DePalma, E. Morant, and C. Breazeal, "Crowdsourcing human-robot interaction: Application from virtual to physical worlds," in *IEEE Int. Sym. on Robots & Human Int. Communications (ROMAN '11)*, 2011, pp. 21–26.
- [10] H. Robbins, "Some Aspects of the Sequential Design of Experiments," *Bulletin of the American Mathematical Society*, vol. 58, pp. 527–535, Sep. 1952.
- [11] P. Auer, N. Cesa-Bianchi, and P. Fischer, "Finite-time analysis of the multiarmed bandit problem," *Machine Learning*, vol. 47, pp. 235–256, 2002.
- [12] J. Current, C. R. Velle, and J. Cohon, "The maximum covering/shortest path problem: A multiobjective network design and routing formulation," *European Journal of Operational Research*, vol. 21, no. 2, pp. 189–199, 1985.
- [13] J. Kuffner, J.J. and S. LaValle, "RRT-connect: an efficient approach to single-query path planning," in *IEEE Int. Conf. on Robotics and Automation (ICRA '00)*, 2000, pp. 995–1001.
- [14] M. Rickert, O. Brock, and A. Knoll, "Balancing exploration and exploitation in motion planning," in *IEEE Int. Conf. on Robotics and Automation (ICRA '08)*, 2008, pp. 2812–2817.
- [15] T. Bektas, "The multiple traveling salesman problem: an overview of formulations and solution procedures," *Omega*, vol. 34, no. 3, pp. 209–219, 2006.
- [16] C. E. Miller, A. W. Tucker, and R. A. Zemlin, "Integer programming formulation of traveling salesman problems," *Journal of the ACM (JACM)*, vol. 7, pp. 326–329, 1960.
- [17] M. R. Garey, R. L. Graham, and D. S. Johnson, "Some NP-complete geometric problems," in *8th annual ACM Symposium on Theory of Computing (STOC '76)*, 1976, pp. 10–22.
- [18] P. Young, "Optimal voting rules," *Journal of Economic Perspectives*, vol. 9, no. 1, pp. 51–64, 1995.