

# Socially-Aware Robot Planning via Bandit Human Feedback

Xusheng Luo<sup>†</sup>, Yan Zhang<sup>†</sup>, Michael M. Zavlanos

Department of Mechanical Engineering and Materials Science, Duke University, USA

{xusheng.luo, yan.zhang2, michael.zavlanos}@duke.edu

**Abstract**—In this paper, we consider the problem of designing collision-free, dynamically feasible, and socially-aware trajectories for robots operating in environments populated by humans. We define trajectories to be social-aware if they do not interfere with humans in any way that causes discomfort. In this paper, discomfort is defined broadly and, depending on specific individuals, it can result from the robot being too close to a human or from interfering with human sight or tasks. Moreover, we assume that human feedback is a bandit feedback indicating a complaint or no complaint on the part of the robot trajectory that interferes with the humans, and it does not reveal any contextual information about the locations of the humans or the reason for a complaint. Finally, we assume that humans can move in the obstacle-free space and, as a result, human utility can change. We formulate this planning problem as an online optimization problem that minimizes the social value of the time-varying robot trajectory, defined by the total number of incurred human complaints. As the human utility is unknown, we employ zeroth order, or derivative-free, optimization methods to solve this problem, which we combine with off-the-shelf motion planners to satisfy the dynamic feasibility and collision-free specifications of the resulting trajectories. To the best of our knowledge, this is a new framework for socially-aware robot planning that is not restricted to avoiding collisions with humans but, instead, focuses on increasing the social value of the robot trajectories using only bandit human feedback.

## I. INTRODUCTION

Recent advances in onboard computing software and hardware make it appear that the time when robots will populate and operate in the same environments with humans is not too far away. As a result, there has recently been a strong interest in developing planning methods for robots that operate in close proximity to humans [1]. To this date, most of these methods focus on safe planning in dense human crowds, where humans move frequently, and they often rely on standard collision avoidance techniques developed in the robotics literature for this purpose. While safety is indeed a critical specification for any human-robot system, the deployment of robots in environments populated by humans also requires planning methods that can incorporate different human utilities that are not always straightforward to express mathematically. In fact, it is often the case that humans will only provide high-level human feedback on whether they like an event or not, e.g., a robot behavior,

without providing relevant contextual information, e.g., their locations in the case of safety. Developing planning methods that can incorporate such unknown human utilities and high-level feedback is necessary to have robots that can efficiently coexist with humans.

In this paper, we consider the problem of planning socially-aware trajectories for robots operating in environments populated by humans. The task assigned to the robot is a basic reach-avoid task. We define a robot trajectory to be social-aware if it does not interfere with humans in any way that causes discomfort. In this paper, discomfort is defined broadly and, depending on specific individuals, it can result from the robot being too close to a human or from interfering with human sight or tasks. Moreover, we assume that human feedback is a bandit feedback indicating a complaint or no complaint on the part of the robot trajectory that interferes with the humans, and it does not reveal any contextual information on the locations of the humans or the reason for a complaint. Finally, we assume that humans can move in the obstacle-free space. As a result, human utilities can change and different robot paths need to be computed for different human configurations. We formulate this planning problem as an online optimization problem that minimizes the social value of the time-varying robot trajectory, defined by the total number of incurred human complaints. Since the human utility function that specifies human feedback is unknown, we can not apply standard gradient-based optimization methods to solve this online optimization problem. Instead, we propose to use derivative-free methods, specifically the zeroth order method, to estimate the gradient by querying humans for feedback on randomly generated trajectories. To obtain a plausible solution, we restrict the number of requests for feedback that can be sent to humans; humans will not provide feedback hundreds of times for the same task. Furthermore, to ensure that the resulting robot trajectories are collision-free and dynamically feasible, we integrate the zeroth order method with an off-the-shelf motion planner that corrects the zeroth-order trajectory to satisfy these additional constraints.

Next, we discuss related work on socially-aware robot navigation, trajectory optimization, and human preference.

### A. Socially-Aware Robot Navigation

To this date, *model-based* and *learning-based* methods are the two main approaches to socially-aware robot naviga-

<sup>†</sup> These Authors contributed equally to this work.

This work is supported in part by AFOSR under award #FA9550-19-1-0169.

tion [2]. Model-based methods typically combine multiagent collision avoidance techniques with predictive models of human motion, such as social force models [3] or gaussian-processes-based statistical models [4]. On the other hand, learning-based methods either learn an interaction model which can predict human motion from a human trajectory dataset by means of hand-crafted features [5], [6], gaussian processes [7], or neural networks [8], or directly enforce multiagent cooperative collision avoidance using deep reinforcement learning [2]. As discussed before, common in these methods is that they all focus on human-robot collision avoidance that can be modeled mathematically. Instead, here we consider socially-aware robot planning for human utilities that are possibly unknown and can not be easily described mathematically. Instead, we only assume that high-level human feedback is available that does not include contextual information about the reasons for a complaint. To the best of our knowledge, this is the first robot planning framework that can incorporate such types of human utilities and feedback.

### B. Trajectory Optimization

Numerical optimization methods for motion planning can be broadly classified into two categories, namely, *derivative-based* and *derivative-free*. Derivative-based methods require a smooth and differentiable objective function and typically solve non-linear motion planning problems by leveraging the first- and second-order derivatives. These methods include, but are not limited to, the well-known motion planners CHOMP [9], TrajOpt [10], KOMO [11], and RieMo [12]. However, derivative-based methods can not handle non-differentiable utility functions, such as the human utilities considered in this work. Such non-differentiable utilities can instead be handled by derivative-free planning methods. A type of derivative-free planning methods that have been widely used in the robotics literature are the sampling-based methods, such as kinodynamic RRT\* [13], which builds graphs by randomly exploring the configuration space and optimizing an objective along the way. Alternatively, population-based optimization methods can be used for motion planning with non-differentiable objectives, such as the particle swarm optimization method [14], that moves a swarm of particles in the search space in search of better robot trajectories. Another popular derivative-free method is STOMP [15] that improves on the current trajectory by applying Monte Carlo simulation to sample a series of noisy trajectories and explore the trajectory space. Although derivative-free planning methods, such as those proposed in [13]–[15], can handle non-differentiable objectives, they still require these objectives to be known and mathematically expressed. This is not the case in the socially-aware planning problem considered here, where human utilities are unknown and only high-level context-unaware human feedback is available. What makes it possible to utilize such human

feedback for robot planning is the proposed zeroth-order optimization framework that is combined with off-the-shelf planners such as the above.

### C. Human Preference Trajectory Optimization

Perhaps most relevant to the method proposed here is recent work on trajectory optimization using human preferences. For example, in [16] task-specific human preferences are incorporated into a motion planner, assuming that these human preferences are formulated as simple convex functions. Similarly, in [17], human preferences are used to iteratively make small improvements to trajectories returned by robots. Instead, here we assume that human feedback is a bandit feedback and the human utility function that governs this feedback is unknown. Following a different approach, [18] adopts a data-driven method to account for human preferences. A set of features regarding the trajectory is designed and a classifier is learned to predict human rating on a new trajectory. Similarly, [19] learns an action policy to tackle grasping tasks by actively querying a human expert for ratings. More broadly, preference-based reinforcement learning [20] focuses on human preference for pairwise comparison between trajectories. Either a policy that satisfies human preference is learned directly [21], [22], or a qualitative preference model capturing a human preference relation or a quantitative utility function estimating a human evaluation criterion are learned and then used to optimize the policy [23]–[26]. While preference-based learning methods can theoretically be used to solve the socially-aware planning problem with bandit human feedback proposed in this paper, their application is not straightforward. First, they can only be applied to stationary problems where the utilities of humans do not change with time. This is because a learned policy for one preference model can not be easily adapted to another preference model corresponding to different human utilities. However, even in the stationary case, the amount of data required by learning-based methods to learn an optimal policy is typically very large. This would make the application of such methods prohibitive in practice, since humans would need to be queried for feedback a prohibitively large number of times. Our proposed framework does not suffer such limitations. Extensive numerical simulations show that it can easily handle changing human utilities and the number of queries for human feedback is low.

The rest of the paper is organized as follows. In Section II, we provide a brief overview of the zeroth order optimization method and in Section III we present the problem formulation. In Section IV we describe our proposed algorithm for the social aware robot planning problem. Numerical experiments are presented in Section V.

## II. PRELIMINARIES

In this section, we provide a short background on zeroth-order optimization since it constitutes the foundation upon

which our proposed method is developed. The zeroth order method [27], [28] is a derivative-free optimization framework that optimizes an objective function based only on the function evaluations, which is a preferred method when the objective function is unknown or noisy, or the gradient is difficult to calculate. Consider an unknown objective function  $g : \mathcal{X} \rightarrow \mathbb{R}$  defined on a  $d$ -dimensional domain  $\mathcal{X} \subset \mathbb{R}^d$ . To minimize the objective  $g$ , the zeroth order method estimates the gradient at the current iterate  $x_k$  by randomly perturbing  $x_k$  and then querying for function evaluations one or multiple times. For example, the following gradient estimate queries for function evaluations twice,

$$\nabla g_x(x_k) \approx \frac{n_x(g(x_k + \delta\mathbf{u}) - g(x_k - \delta\mathbf{u}))}{2\delta}\mathbf{u}, \quad (1)$$

where  $n_x$  is the size of the decision variable  $x$ ,  $\mathbf{u}$  is a random vector sampled from a uniform distribution in a unit sphere or a Gaussian distribution in  $\mathbb{R}^{n_x}$  and  $\delta > 0$  is the exploration parameter. As discussed in [27], [28], the expression in (1) is an unbiased stochastic gradient estimate of the smoothed function  $\mathbb{E}_{\mathbf{u}}[g(x_k + \delta\mathbf{u})]$ . Therefore, when the zeroth order method updates the current iterate  $x_k$  using (1), it in fact minimizes the smoothed objective function  $\mathbb{E}_{\mathbf{u}}[g(x_k + \delta\mathbf{u})]$ , whose distance to  $g(x_k)$  is bounded by an amount proportional to the exploration parameter  $\delta$ , see [27], [28]. After obtaining the estimated gradient, the next iterate is computed in a usual way as

$$x_{k+1} = x_k - \eta \nabla g_x(x_k), \quad (2)$$

where  $\eta > 0$  is the step size.

### III. PROBLEM DEFINITION

Consider an environment  $\mathcal{W} \subset \mathbb{R}^w, w \in \{2, 3\}$ , let  $\mathcal{O} = \{o_1, \dots, o_m\}$  denote the set of  $m$  obstacles in  $\mathcal{W}$ , and assume that the obstacle-free space  $\mathcal{W} \setminus \mathcal{O}$  is populated by humans that are distributed possibly in a spatially uneven way. Assume also that the model of human dynamics is unknown and human movement is random, making it unlikely to predict the moving pattern. Given a group of  $p$  humans in the environment, let  $\mathcal{L}_{T_i} = \{l_{1,T_i}, \dots, l_{p,T_i}\}$  be the set that collects their locations at discrete time instants  $T_i, i \in \{0, \dots, K\}$ . We make the following assumption on the motion pattern of the humans.

*Assumption 3.1:* There exists a constant  $\delta T > 0$  such that  $T_{i+1} - T_i > \delta T$  for all  $i = 0, \dots, K - 1$ . Moreover, during every time instant  $T_{i+1}$ , humans  $k \in \{1, \dots, p\}$  move instantaneously between consecutive locations  $l_{k,T_i}$  and  $l_{k,T_{i+1}}$ , and they remain stationary during the time interval  $(T_i, T_{i+1})$ .

The existence of the bound  $\delta T > 0$  in Assumption 3.1 ensures that the times when humans move in the environment are not arbitrarily close to each other. Since humans also move instantaneously, Assumption 3.1 captures situations

where human motion is infrequent and sporadic. For example, it models typical work environments, where humans remain around certain locations for a while and then quickly move to other locations to perform subsequent tasks.

Consider also a robot navigating in the obstacle-free environment  $\mathcal{W} \setminus \mathcal{O}$  from a starting location  $x_0$  to a target location  $x_g$ . Let  $q \in \mathcal{Q}$  denote the state of the robot that includes the robot's location, orientation, etc, where  $\mathcal{Q}$  is the feasible state space. Let also  $u \in \mathcal{U}$  denote the control input of the robot, where  $\mathcal{U}$  is the set of admissible control inputs. Moreover, let  $t_j \in [T_i, T_{i+1})$  for  $j \in \{0, 1, \dots, n\}$ , denote a sequence of time instants defined over the time intervals that humans remain stationary, and define the discrete-time dynamics of the robots during this time interval as

$$q_{t_{j+1}} = f(q_{t_j}, u_{t_j}). \quad (3)$$

Finally, define a discrete robot trajectory during the time interval  $[T_i, T_{i+1})$  by the sequence of robot positions  $\mathbf{x}_{T_i} = x_{t_0}, x_{t_1}, \dots, x_{t_n}$ , where  $x_{t_j} \in \mathcal{W} \setminus \mathcal{O}$  is the position component of  $q_{t_j}$  and  $t_j \in [T_i, T_{i+1})$  for all  $j \in \{0, \dots, n\}$ . Particularly,  $x_{t_0} = x_0$  is the starting location and  $x_{t_n} = x_g$  is the target location. In what follows, with slight abuse of notation, we write  $\mathbf{x}_{T_i} = [x_{t_0}^\top, x_{t_1}^\top, \dots, x_{t_n}^\top]^\top$ .

Our goal is to design a collision-free, dynamically feasible, and socially-aware trajectory for the robot that does not interfere with humans in any way that causes discomfort. In this paper, discomfort is defined broadly and, depending on specific individuals, it can result from the robot being too close to a human, or from interfering with human sight or tasks. To do so, in what follows, we propose a novel way to incorporate human feedback in the robot planning process. Specifically, we make the following assumptions on the feedback provided by humans.

*Assumption 3.2:* Human feedback is a bandit feedback indicating a complaint or no complaint on the part of the robot trajectory that falls in human vicinity. Moreover, human feedback does not reveal any contextual information about the locations of the humans or the reason for a complaint. Finally, the human utility function does not change during time intervals  $(T_i, T_{i+1})$ , but it can change at time instants  $T_i$  when humans move.

We model human feedback using an implicit human-and-time-dependent utility function  $h_{k,T} : \mathbf{x}_T \mapsto \{0, 1\}$  which indicates the preference of individual  $k$  regarding the trajectory  $\mathbf{x}_T$  at time instant  $T$ , where we drop the subscript  $i$  from  $T_i$  when the notation is clear. Particularly, the function  $h_{k,T}$  returns 1 if individual  $k$  complains about the trajectory at time instant  $T$ , otherwise it returns 0. Dependence of  $h_{k,T}$  on time  $T$  captures the fact that an individual's feedback on a trajectory may vary over time. For instance, humans may become more or less tolerant to a robot moving in their vicinity, depending on the type of their ongoing work. Essentially, we treat each individual as a black box and respect their preferential heterogeneity by means of distinct

functions  $h_{k,T}$ . We capture the *social value* of the trajectory  $\mathbf{x}_T$  by the total number of complaints incurred by that trajectory, defined as  $h_T(\mathbf{x}_T) = \sum_{k=0}^p h_{k,T}(\mathbf{x}_T)$ . Then, a trajectory  $\mathbf{x}_T$  is said to be *socially-aware* if it incurs zero complaints, i.e.,  $h_T(\mathbf{x}_T) = 0$ . In what follows we avoid relying on models of human feedback to achieve socially-aware planning [16], [18], [25], since such models can differ across humans and vary over time.

On the other hand, we define a collision between the trajectory  $\mathbf{x}_T$  and an obstacle  $o_j$  if a line segment between consecutive waypoints in  $\mathbf{x}_T$  intersects the obstacle  $o_j$ . Then the trajectory  $\mathbf{x}_T$  is *collision-free* if it never collides with any obstacle.<sup>1</sup> Let  $\mathcal{C}$  be a set that collects all collision-free discrete trajectories in the environment  $\mathcal{W} \setminus \mathcal{O}$ . Finally, we define a *dynamically-feasible* trajectory to be one that respects the dynamical robot constraints (3), and collect all dynamically-feasible discrete trajectories in the environment  $\mathcal{W} \setminus \mathcal{O}$  in the set  $\mathcal{D}$ . Then, the socially-aware robot planning problem considered in this paper can be formulated as

$$\begin{aligned} \min_{\mathbf{x}_{T_0}, \dots, \mathbf{x}_{T_{K-1}}} \quad & \sum_{i=0}^{K-1} h_{T_i}(\mathbf{x}_{T_i}) \\ \text{s.t.} \quad & \mathbf{x}_{T_i} \in \mathcal{C} \cap \mathcal{D}, \forall i \in \{0, \dots, K-1\}. \end{aligned} \quad (4)$$

Problem (4) is an *offline* optimization problem and it can be solved after time instant  $T_{K-1}$  assuming that all utility functions  $h_{T_i}$  are known in hindsight. However, if the human utilities are unknown and a new socially-aware trajectory has to be designed during every time interval  $[T_i, T_{i+1})$ , problem (4) needs to be solved in an *online* fashion. We can do this by solving the following optimization problem during every time interval  $[T_i, T_{i+1})$ ,

$$\min_{\mathbf{x}_{T_i}} h_{T_i}(\mathbf{x}_{T_i}) \quad \text{s.t. } \mathbf{x}_{T_i} \in \mathcal{C} \cap \mathcal{D} \quad (5)$$

Note that the objectives in problem (5) are time-varying due to the varying human locations and feedback. Based on problem (5), the socially-aware robot planning problem considered in this paper can be formulated as follows.

*Problem 1:* Let Assumptions 3.1 and 3.2 hold. Given a pair of starting and target locations and a sequence of discrete time instants  $T_0, T_1, \dots, T_K$ , with  $K \in \mathbb{N}$ , design a sequence of collision-free, dynamically-feasible, and socially-aware trajectories  $\mathbf{x}_{T_0}, \mathbf{x}_{T_1}, \dots, \mathbf{x}_{T_{K-1}}$  each of which solves the optimization problem (5).

The solution of an online optimization problem should track the solution of its offline counterpart [29]. As such, we can measure the quality of the sequence of trajectories returned by the solution of Problem 1 using the *cumulative regret* over the whole time period defined as

$$R_K = \sum_{i=0}^{K-1} h_{T_i}(\tilde{\mathbf{x}}_{T_i}) - \sum_{i=0}^{K-1} h_{T_i}(\mathbf{x}_{T_i}^*), \quad (6)$$

<sup>1</sup>Note that, in this paper, collision avoidance is defined in terms of obstacles. Collisions with humans are captured by human complaints.

where  $\tilde{\mathbf{x}}_{T_0}, \dots, \tilde{\mathbf{x}}_{T_{K-1}}$  is the computed solution sequence of problem (5) and  $\mathbf{x}_{T_0}^*, \dots, \mathbf{x}_{T_{K-1}}^*$  is the minimizer sequence of problem (4) that incurs the least number of complaints if the functions  $h_{T_i}$  are known in hindsight. The cumulative regret is a dynamic regret since the optimal trajectory  $\mathbf{x}_{T_i}^*$  varies due to human motion and the variations in the human feedback function  $h_{T_i}$  over time. As shown in [30], the best known bound for the dynamic regret (6) grows at the same rate at which the optimal trajectory  $\mathbf{x}_{T_i}^*$  varies over time. Therefore, the cumulative utility (6) shall at best grow at the same rate at which the human positions and preferences change. Note that since the human utility  $h_{T_i}(\mathbf{x}_{T_i}^*)$  is unknown, we can not employ standard gradient-based optimization methods to solve Problem 1. Instead, we employ the zeroth order optimization, discussed in Section II, to solve this online problem.

#### IV. SOCIALLY-AWARE ROBOT PLANNING

In this section, we first reformulate Problem 1 by manipulating the constraint set in (5) and then develop a novel method that combines zeroth order optimization with Model Predictive Control (MPC) [31]–[33] to solve the socially-aware robot planning Problem 1. Specifically, let  $\mathcal{F}(\mathbf{x})$  be the set of collision-free and dynamically-feasible trajectories that are in the neighborhood of a reference trajectory  $\mathbf{x}$  and define the problem

$$\min_{\mathbf{x}_{T_i}} h_{T_i}(\mathbf{x}_{T_i}) \quad \text{s.t. } \mathbf{x}_{T_i} \in \mathcal{F}(\mathbf{x}_{T_i}). \quad (7)$$

Note that problems (5) and (7) are equivalent. To see this, observe that if  $\mathbf{x}_{T_i}^*$  is the optimal solution of (5) then it is collision-free and dynamically feasible. Therefore, it belongs to the set  $\mathcal{F}(\mathbf{x}_{T_i}^*)$ . Since (5) and (7) share the same cost function, this means that  $\mathbf{x}_{T_i}^*$  is also the optimal solution of (7). On the other hand, if  $\mathbf{x}_{T_i}^*$  is the optimal solution of (7) then it is collision-free and dynamically feasible. Therefore, it satisfies the constraints of problem (5), which means that  $\mathbf{x}_{T_i}^*$  is also the optimal solution of (5).

Computing the set  $\mathcal{F}(\mathbf{x})$  is generally difficult. Instead, we can compute a single trajectory in  $\mathcal{F}(\mathbf{x})$  using off-the-shelf motion planners, such as MPC. Let  $m(\mathbf{x})$  denote such a trajectory, which satisfies  $m(\mathbf{x}) \in \mathcal{F}(\mathbf{x})$ . Then, a sufficient condition for feasibility of problem (7) is that  $\mathbf{x}_{T_i} = m(\mathbf{x}_{T_i})$ . To enforce this constraint, we consider  $m(\mathbf{x}_{T_i})$  as a target trajectory that  $\mathbf{x}_{T_i}$  needs to track. Moreover, we query humans for feedback on the trajectory  $m(\mathbf{x}_{T_i})$  rather than the trajectory  $\mathbf{x}_{T_i}$ . This is because  $\mathbf{x}_{T_i}$  may not be collision-free or dynamically feasible, while  $m(\mathbf{x}_{T_i})$  is. Therefore, it is not meaningful to request human feedback on an infeasible trajectory. With the above modifications, we obtain the following penalty-based reformulation of problem (7) as

$$\min_{\mathbf{x}_{T_i}} h_{T_i}(m(\mathbf{x}_{T_i})) + \rho \|\mathbf{x}_{T_i} - m(\mathbf{x}_{T_i})\|, \quad (8)$$

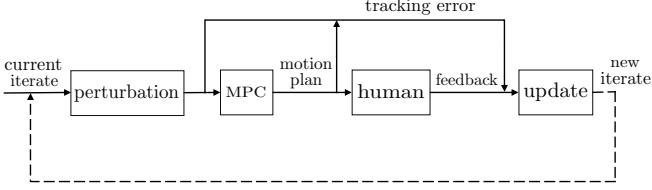


Figure 1. Zeroth order method for socially-aware robot planning

where  $\rho$  is a positive penalty parameter and  $\|\mathbf{x}_{T_i} - m(\mathbf{x}_{T_i})\|$  measures the tracking error between  $\mathbf{x}_{T_i}$  and the motion planner. Note that (8) is not necessarily equivalent to (7). This is because there is no guarantee that a motion planner can return a trajectory  $m(\mathbf{x}_{T_i}^*)$  that satisfies  $\mathbf{x}_{T_i}^* = m(\mathbf{x}_{T_i}^*) \in \mathcal{F}(\mathbf{x}_{T_i}^*)$ , where  $\mathbf{x}_{T_i}^*$  is the optimal solution of (7). Note also that the solution of (8) returns a trajectory  $\mathbf{x}_{T_i}$  but the robot executes the trajectory  $m(\mathbf{x}_{T_i})$  that is guaranteed to be collision-free and dynamically feasible. Finally, observe that in (8), the goal is to minimize human complaints, but other objectives such as control effort can also be incorporated in this formulation. For example, the MPC trajectory  $m(\mathbf{x}_{T_i})$  in (8) can be easily constructed in a way that minimizes such additional objectives; see, e.g., Section IV-B..

As discussed in Section III, since the human utility function  $h_T$  is unknown, we need to utilize the zeroth order method to solve problem (8), which we combine with an MPC motion planner to obtain the trajectories  $m(\mathbf{x}_{T_i})$  in (8) as follows. At every iteration of the proposed zeroth order method, the current trajectory iterate is perturbed and then rectified by the MPC motion planner so that it is collision-free and dynamically feasible. Then, the humans are queried for feedback on the perturbed trajectories and this feedback along with the tracking error is used by the zeroth order method to estimate the gradient and update the current trajectory iterate; see Section II. The proposed algorithm is illustrated in Figure 1. In what follows, we discuss the proposed zeroth order method, the MPC motion planning problem, and the integrated system in detail.

#### A. Zeroth Order Optimization Method

The advantage of the proposed zeroth-order method is that it can be used to incorporate bandit human feedback in motion planning. Since only the value of the human utility function  $h_T(m(\mathbf{x}_T))$  is accessible, the proposed zeroth order method uses these values to estimate the gradient of human feedback as

$$\nabla h_T(m(\mathbf{x}_T^k)) \approx \frac{n_{\mathbf{x}} (h_T(m(\mathbf{x}_T^{k,+})) - h_T(m(\mathbf{x}_T^{k,-})))}{2\delta} \mathbf{u}, \quad (9)$$

where  $\mathbf{x}_T^{k,\pm} = \mathbf{x}_T^k \pm \delta \mathbf{u}$ ,  $\mathbf{x}_T^k$  is the  $k$ -iterate of  $\mathbf{x}_T$ ,  $m(\mathbf{x}_T)$  is the trajectory returned by the MPC motion planner that will be introduced in the next section,  $n_{\mathbf{x}}$  is the size of vector  $\mathbf{x}_T$ , and  $\mathbf{u}$  is a random vector of unit length. Given

the vector  $\mathbf{u}$ , the trajectories  $m(\mathbf{x}_T^k + \delta \mathbf{u})$  and  $m(\mathbf{x}_T^k - \delta \mathbf{u})$  are deterministic.

A main limitation of zeroth-order methods in general is that their convergence is slow, especially when the problem dimension is high [27], [28]. This is the case here as the number of waypoints  $n_{\mathbf{x}}$  in the trajectory  $\mathbf{x}_T$  can be very large. To address this limitation, we propose local updates of the trajectory  $\mathbf{x}_T$ , restricted to those waypoints that are directly affected by human complaints. This is a reasonable assumption, since humans are more likely to complain on that part of the trajectory that causes them discomfort rather than on the whole robot trajectory, which may not even be fully observable. To develop the proposed local update scheme, we replace the utility functions  $h_{k,T}(m(\mathbf{x}_T))$  by functions  $h_{k,T}([m(\mathbf{x}_T)]_{\mathcal{S}^{k,T}})$ , where the index set  $\mathcal{S}^{k,T}$  captures the local awareness of human  $k$ , that is, those waypoints in the whole trajectory  $m(\mathbf{x}_T)$  that directly affect human  $k$  at time instant  $T$ . In this case, to minimize the utility function  $h_T(m(\mathbf{x}_T)) = \sum_{k=0}^p h_{k,T}([m(\mathbf{x}_T)]_{\mathcal{S}^{k,T}})$ , it is sufficient to only update the waypoints in the set  $\mathcal{S}_p \triangleq \cup_k \mathcal{S}^{k,T}$  according to the modified zeroth order update

$$\nabla h_T^{\mathcal{S}_p}(m(\mathbf{x}_T^k)) \approx \frac{w |\mathcal{S}_p| (h_T(m(\mathbf{x}_T^{k,+})) - h_T(m(\mathbf{x}_T^{k,-})))}{2\delta} \mathbf{u}^{\mathcal{S}_p}, \quad (10)$$

where  $\mathbf{x}_T^{k,\pm} = \mathbf{x}_T^k \pm \delta \mathbf{u}^{\mathcal{S}_p}$ ,  $w \in \{2, 3\}$  is the dimensionality of a waypoint,  $|\mathcal{S}_p|$  is the cardinality of the set  $\mathcal{S}_p$  and the vector  $\mathbf{u}^{\mathcal{S}_p}$  is computed by setting the entries of  $\mathbf{u}$  that are not in the set  $\mathcal{S}_p$  to be 0 and then normalizing, i.e.,  $\mathbf{u}^{\mathcal{S}_p} = \mathbf{u}^{\mathcal{S}_p} / \|\mathbf{u}^{\mathcal{S}_p}\|$ . Since  $|\mathcal{S}_p| \ll n_{\mathbf{x}}$ , the convergence speed of the algorithm can be greatly improved. Similarly, the gradient of the tracking error between the zeroth-order iterate and the MPC motion plan can be obtained as

$$\nabla e^{\mathcal{S}_p}(\mathbf{x}_T^k) \approx \frac{w |\mathcal{S}_p| (e(\mathbf{x}_T^{k,+}) - e(\mathbf{x}_T^{k,-}))}{2\delta} \mathbf{u}^{\mathcal{S}_p}, \quad (11)$$

where  $e(\mathbf{x}_T^k) = \|\mathbf{x}_T^k - m(\mathbf{x}_T^k)\|$  denotes the tracking error. Note that  $\nabla e^{\mathcal{S}_p}(\mathbf{x}_T^k)$  can also be computed using finite differences since this objective is known. In what follows, we discuss how to compute the set  $\mathcal{S}_p$ .

Since each trajectory needs to connect the starting and target locations of the robot, we assume that the first and last waypoints of the trajectory  $\mathbf{x}_T$  do not affect any human and, therefore, they are not included in the set  $\mathcal{S}_p$ . All other locations in the free space, if they are traversed by the trajectory  $\mathbf{x}_T$ , can possibly cause discomfort to humans. In what follows, we discuss two forms of human feedback, thus, resulting in two ways to decide the set  $\mathcal{S}_p$ . (i) Humans can only send complaints without revealing which waypoints affect them. In this case, the set  $\mathcal{S}_p$  is the full trajectory and Alg. 1 perturbs the whole trajectory  $\mathbf{x}_T$ . We refer to this as the full perturbation scheme. (ii) Humans send complaints and also report the part of the trajectory that affects them.

In this case, the set  $\mathcal{S}_p$  contains those waypoints indicated by the human as well as a random number of waypoints before and after the reported sequence of waypoints. This provides more room to the zeroth-order method to compute trajectories that can be easier modified by the motion planner. Moreover, these additional waypoints can account for human errors in reporting the waypoints that cause them discomfort. We refer to this as the local perturbation scheme.

### B. Model Predictive Control (MPC)

MPC has long been used to plan collision-free robot trajectories that track predefined reference paths [34]–[36]. We adopt the approach proposed in [35]. Specifically, given a reference trajectory  $x_{t_0}, x_{t_1}, \dots, x_{t_n}$  returned by the zeroth order method that can be dynamically infeasible or can collide with obstacles, our goal is to compute the optimal control sequence  $\{u_{t_j,k}^*\}_{k=0}^{N_m-1}$ , at each time instant  $t_j, j \in \{0, \dots, n-1\}$ , by solving the following optimal control problem over a horizon of length  $N_m$ :

$$\begin{aligned} \{u_{t_j,k}^*\}_{k=0}^{N_m-1} &= \operatorname{argmin} J(x_{t_j}, \{u_{t_j,k}\}_{k=0}^{N_m-1}) \\ &= \mu \sum_{k=0}^{N_m-1} J^{\text{obs}}(x_{t_j,k}, x_{t_j,k+1}) + \frac{1}{2} \tilde{x}_{t_j,N_m}^\top P \tilde{x}_{t_j,N_m} \\ &+ \sum_{k=0}^{N_m-1} \left( \frac{1}{2} \tilde{x}_{t_j,k}^\top Q \tilde{x}_{t_j,k} + \frac{1}{2} u_{t_j,k}^\top R u_{t_j,k} \right) \end{aligned} \quad (12a)$$

$$\text{s.t. } q_{t_j,k+1} = f(q_{t_j,k}, u_{t_j,k}) \quad (12b)$$

$$u_{t_j,k} \in \mathcal{U}, \forall k \in \{0, 1, \dots, N_m - 1\} \quad (12c)$$

where  $\tilde{x}_{t_j,k} = x_{t_j+k} - x_{t_j,k}$  is the deviation from the reference trajectory, and  $\mu, P, Q, R$  are constant weighting matrices. The first term in (12a) penalizes collisions with obstacles, the second term penalizes the terminal deviation, and the last term penalizes the deviation from the reference trajectory and the control input. Note that in the proposed MPC problem the penalty on the control input can help reduce the trajectory length. Moreover, the penalty term  $J^{\text{obs}}(x_{t_j,k}, x_{t_j,k+1})$  on the collisions is a point-wise potential function, defined as

$$J^{\text{obs}}(x_{t_j,k}, x_{t_j,k+1}) = \frac{1}{\min_i d(o_i, \overline{x_{t_j,k} x_{t_j,k+1}}) + \epsilon}, \quad (13)$$

where  $\min_i d(o_i, \overline{x_{t_j,k} x_{t_j,k+1}})$  is the shortest distance of the line segment that connects waypoints  $x_{t_j,k}$  and  $x_{t_j,k+1}$ , denoted by  $\overline{x_{t_j,k} x_{t_j,k+1}}$ , to any obstacle,  $\epsilon$  is a small positive constant to avoid the singularity.

### C. The Integrated Socially-Aware Planning Algorithm

Alg. 1 shows the proposed online solution to Problem 1 during the whole time period  $T_0, T_1, \dots, T_K$ , that integrates zeroth-order optimization for social awareness of the desired trajectory and MPC for collision avoidance and dynamic feasibility. As shown in [37], if the objective function is

non-degenerate, which is a weaker condition than the strong convexity, the dynamic regret in (6) can be improved by estimating the gradient multiple times per time instant. Although, compared to [37], here the form of  $h_{T_i}$  is unknown, we adopt the same idea and run  $N$  gradient updates at every time instant. Note that the inner iteration [line 2] of Alg. 1, reduces to a common online algorithm when  $N = 1$ . Within each inner iteration, the humans are queried for feedback twice: one for the two perturbed trajectories and one for the evaluation of the updated trajectory. Therefore, the number of queries during each time interval  $[T_i, T_{i+1})$  is twice as much as the number of inner iterations. Note that the number of queries can be further reduced by running the  $N$  inner iterations without checking whether the updated trajectory has 0-complaints or not. In this case, humans are only queried once per iteration. In what follows, we discuss Alg. 1 in details.

A good initial trajectory  $\mathbf{x}_{T_0}^0$ , starting from the starting location  $x_0$  and ending at the target location  $x_g$ , can be a collision-free trajectory that is provided by motion planning algorithms, such as RRT\* [38]. But it can also be any other trajectory without considering collision since MPC will address this. For each time interval  $[T_i, T_{i+1})$ ,  $\forall i \in \{0, \dots, K-1\}$ , Alg. 1 runs multiple iterations to obtain a trajectory with high social awareness. Specifically, at every iteration  $k$ , Alg. 1 first collects the human feedback on the trajectory  $m(\mathbf{x}_{T_i}^k)$ . If  $m(\mathbf{x}_{T_i}^k)$  has zero complaints or a predefined maximum number of iterations is reached, then the inner iteration terminates and the trajectory returned by MPC is the solution for the time interval  $[T_i, T_{i+1})$  [lines 3-5]. Otherwise, Alg. 1 updates the waypoints in  $\mathcal{S}_p$ , perturbs the current iterate  $\mathbf{x}_{T_i}^k$  by adding and subtracting  $\delta \mathbf{u}^{\mathcal{S}_p}$ , and queries the humans for feedback twice [lines 6-7]. Then, the algorithm computes the gradient for the objective function and updates the iterate [lines 8-9]. We add an extra weight parameter  $\alpha > 0$  before the gradient of the utility function. At the next time instant  $T_{i+1}$ , the humans in the environment may move. The trajectory from the previous instant  $T_i$  is used as a warm start to initialize  $\mathbf{x}_{T_{i+1}}^0$  at instant  $T_{i+1}$ .

## V. NUMERICAL EXPERIMENTS

In this section, we present multiple case studies, implemented using Python 3.6.3 on a computer with a 2.3 GHz Intel Core i5 processor and 8G RAM, that illustrate the efficiency of the proposed algorithm. We consider a square continuous environment  $\mathcal{W}$  of size  $20 \times 20$  where two square obstacles of size  $2 \times 2$  are located near the center. The starting location of the robot is at  $(0, 0)$  and the target location is at  $(20, 20)$ ; see also Fig. 2. The start and target locations are considered fixed throughout the simulations. We generate a trajectory  $\mathbf{x}_0$  that consists of 15 waypoints that connects the initial and target locations; see red arrow path in Fig. 2. This trajectory constitutes the initial trajectory  $\mathbf{x}_{T_0}^0$  for the proposed zeroth-order method.

**Algorithm 1:** Online trajectory planning based on human feedback

---

**Input:** Initial trajectory  $\mathbf{x}_{T_0}^0$ , number of time instants  $K$ , maximum number of iterations  $N$ , weights  $\alpha, \rho$ , exploration parameter  $\delta$ , step size  $\eta$ .

- 1 **for**  $0 \leq i \leq K - 1$  **do**
- 2   **for**  $0 \leq k \leq N$  **do**
- 3     **if**  $h_{T_i}(m(\mathbf{x}_{T_i}^k)) = 0$  or  $k = N$  **then**
- 4        $\mathbf{x}_{T_i}^k = m(\mathbf{x}_{T_i}^k)$ ;
- 5       **break** ;
- 6     Get the set  $\mathcal{S}_p$  of waypoints to perturb ;
- 7     Get the human feedback  $h_T(m(\mathbf{x}_{T_i}^k + \delta \mathbf{u}_p^{\mathcal{S}}))$  and  $h_T(m(\mathbf{x}_{T_i}^k - \delta \mathbf{u}_p^{\mathcal{S}}))$  ;
- 8     Estimate gradient:
- 9        $\nabla f_{T_i}(\mathbf{x}_{T_i}^k) \approx \alpha \nabla h_{T_i}^{\mathcal{S}_p}(m(\mathbf{x}_{T_i}^k)) + \rho \nabla e^{\mathcal{S}_p}(\mathbf{x}_{T_i}^k)$ .
- 10    Update iterate:
- $\mathbf{x}_{T_i}^{k+1} = \mathbf{x}_{T_i}^k - \eta \nabla f_{T_i}(\mathbf{x}_{T_i}^k)$ .
- 10     $\mathbf{x}_{T_{i+1}}^0 = \mathbf{x}_{T_i}$ ;

---

Furthermore, we consider a group of  $p \in \{20, 30, 40, 50, 60\}$  humans in the environment  $\mathcal{W} \setminus \mathcal{O}$ . At time  $T_0$ , we assume that the human population is randomly located in the environment. Then, at each time instant  $T_i$ , each individual moves by uniformly sampling a location within a circle of radius of  $r \in \{0.3, 0.5, 1\}$  centered at their current location. The size of the circle captures the range of an individual's movement. To quantify human feedback, we assume that each individual possesses a restricted zone and makes a complaint if there exists a trajectory segment passing through the restricted zone. The restricted zone is represented as a circle of radius  $r_k$ , where  $r_k$  is human-dependent and it reflects the sensitivity of an individual to violation of their restricted zone; see also Fig. 2. For every human  $k \in \{1, \dots, p\}$ , we let  $r_k$  randomly take a value in the set  $\{0.3, 0.4, 0.5, 0.7\}$ . The radius of the restricted zone is randomly updated at subsequent time instants  $T_i$ , which captures the temporal variation of a human's feedback. We emphasize that the restricted zone is only a construct used to generate complaints and it is not known by the zeroth order planner; neither are the human locations. Finally, we consider a time period containing the sequence of time instants  $T_0, \dots, T_{30}$ . The parameters of the zeroth order method in Alg. 1 are set as  $\alpha = 10, \rho = 1, \delta = 10$  and  $\eta = 0.1$  for the full perturbation scheme and  $\eta = 0.5$  for the local perturbation scheme.

#### A. Robot Model

We consider a unicycle robot model [39]. Let  $x = (x^1, x^2) \in \mathcal{W} \subset \mathbb{R}^2$  and  $\theta \in (-\pi, \pi]$ , respectively, represent

Table I  
STATISTICS FOR VARYING NUMBER OF HUMANS

$p$	$N_{\text{full}}$	$N_{\text{local}}$	$d_{\text{full}}$	$d_{\text{local}}$
20	$2.35 \pm 2.37$	$3.25 \pm 4.44$	$31.88 \pm 2.42$	$30.34 \pm 0.99$
30	$5.58 \pm 5.19(1)$	$5.20 \pm 8.33$	$34.50 \pm 4.70$	$31.46 \pm 2.84$
40	$12.20 \pm 7.67$	$7.75 \pm 4.63$	$39.36 \pm 6.14$	$35.57 \pm 6.51$
50	$12.35 \pm 13.10$	$7.60 \pm 8.11$	$39.56 \pm 8.52$	$35.60 \pm 6.18$
60	$18.33 \pm 13.17(5)$	$8.68 \pm 8.52(1)$	$43.92 \pm 8.06$	$38.01 \pm 8.23$

$N_{\text{full}}$  and  $N_{\text{local}}$  represent the number of iterations taken by the full and local perturbation schemes, respectively, and  $d_{\text{full}}$  and  $d_{\text{local}}$  represent the trajectory length returned by full and local perturbation schemes, respectively. The numbers in the parentheses next to certain results report the number of trials out of 20 where no socially-aware trajectory is generated within 50 iterations. During these trials, paths with high social value were still generated, but the number of complaints was not zero.

the location and orientation of the robot, and  $u = (v, \omega) \in \mathcal{U}$  denote the control input, where  $v$  and  $\omega$  are the linear and angular velocity, respectively. Then the discrete kinematic equation governing the state  $q = (x, \theta)$  is

$$q_{t_{j+1}} = f(q_{t_j}, u_{t_j}) = q_{t_j} + \Delta t \begin{bmatrix} \cos \theta_{t_j} & 0 \\ \sin \theta_{t_j} & 0 \\ 0 & 1 \end{bmatrix} u_{t_j}, \quad (14)$$

where  $\Delta t = t_{j+1} - t_j$ . In the simulation experiments that follow, the MPC parameters are set as  $N_m = 5, \mu = 50, P = Q = \text{diag}([25, 25]), R = \text{diag}([10, 1]), \epsilon = 10^{-8}, v \in [0.1, 5], \omega \in (-\pi, \pi], \Delta t = 1$ .

#### B. Socially-Aware Planning in Stationary Environments

We first validate Alg. 1 and compare the full perturbation and local perturbation methods, discussed in Section IV-A, in situations where humans do not move. We set  $K = 1$  (consider  $T_0$  only) and  $N = 50$  in Alg. 1 and, as discussed in Section IV-C, we perform up to  $N$  gradient updates until the inner iteration terminates, with querying humans for feedback on the updated trajectory, so that the total number of queries is twice as much as the number of iterations that Alg. 1 takes. Moreover, we use the trajectory  $\mathbf{x}_0$  as the initial trajectory  $\mathbf{x}_{T_0}^0$  for both full perturbation and local perturbation schemes. The number  $p$  of humans varies from 20 to 60. For a specific value of  $p$ , we report the mean and standard deviation of the number of iterations taken to find a collision-free, dynamically-feasible and socially-aware trajectory, averaged over 20 trials; see also Table I. We also present trajectory length as another metric to compare the full and local perturbation methods. Observe that, for each scheme, both the average number of iterations required to obtain a socially-aware trajectory and the trajectory length grow as  $p$  increases. In comparison, the local perturbation scheme requires fewer iterations and generates shorter-length trajectories than the full perturbation method. Note that the average number of queries for human feedback, of Alg. 1 is low; between 20 and 40 times for the local and full perturbation methods, respectively. Therefore, the proposed method can be used in practice without burdening humans

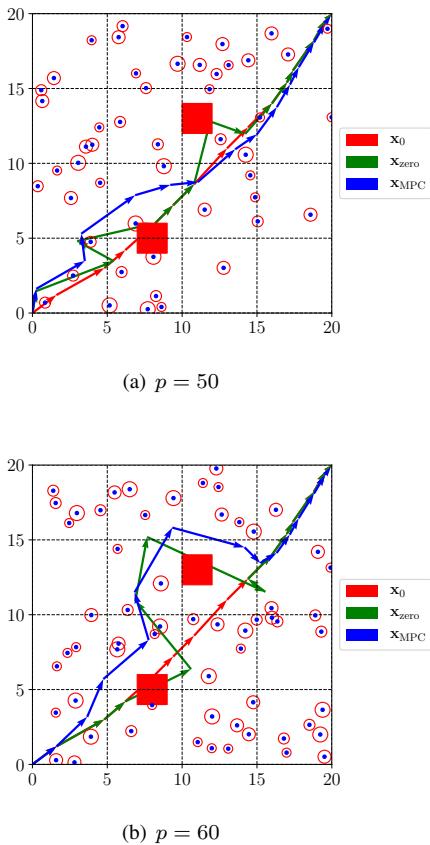


Figure 2. Sample trajectories. The blue dots represent humans and the red circles are restricted zone. The red squares represent obstacles. The red arrow path is the initial trajectory, and the green and blue arrow path are returned by the zeroth order method and MPC motion planner, respectively.

with prohibitively many feedback requests. In contrast, the reinforcement learning-based methods discussed in Section I-C will require orders of magnitude more samples to obtain an optimal policy. Two sample trajectories returned by the local perturbation scheme, when  $p = 50, 60$ , are shown in Fig. 2. We see that the trajectory returned by the MPC motion planner successfully avoids human restricted zone and obstacles, whereas its reference trajectory returned by the zeroth order method violates the restricted zone and collides with obstacles.

### C. Socially-Aware Planning in Dynamic Environments

In this section, we consider situations where humans move at time instants  $T_i, i \in \{0, \dots, 30\}$  and their corresponding utilities change. Therefore, we solve the online problem (8). At each time instant, the full and local perturbation schemes are used for the same human distribution and human utility function. We assume that there always exists a dynamically-feasible, collision-free and socially-aware trajectory during each time interval. That is, each trajectory in the minimizer sequence  $\mathbf{x}_{T_0}^*, \dots, \mathbf{x}_{T_{K-1}}^*$  of problem (4) incurs zero com-

plaints. In this case, the cumulative regret in (6) is reduced to

$$R_K = \sum_{i=0}^{K-1} h_{T_i}(\tilde{\mathbf{x}}_{T_i}) \quad (15)$$

In the following simulation, we fix the number of humans at 50 and vary the movement radius  $r \in \{0.3, 0.5, 1\}$  and the maximum number  $N \in \{1, 2, 3, 4\}$  of gradient updates in Alg. 1. Different from V-B, we perform the  $N$  gradient updates without querying humans for feedback so that the total number of queries is equal to the number of iterations of Alg. 1. The value or the radius  $r$  affects the relative change between human distributions and, therefore, the relative change between optimization problems at consecutive time intervals. For each combination of  $r$  and  $N$ , we run 5 trials and report the mean cumulative regret as a function of the time instants  $T_i$ , for the full and local perturbation schemes, respectively. As a baseline, we also report the mean cumulative regret incurred by the initial trajectory  $m(\mathbf{x}_{T_0})$  returned by MPC, at every time interval  $[T_i, T_{i+1})$  assuming that this is also the trajectory returned by the algorithm at that time interval. This baseline is equivalent to the case where  $N = 0$ . Fig. 3 depicts the growth of the regret  $R_T$  with respect to the maximum number of iterations  $N$  for different human motion radii  $r$ . Observe that, compared to the baseline trajectory  $m(\mathbf{x}_{T_0})$ , the full and local perturbation schemes can improve the social value of the trajectory with few queries. Particularly, using only 3 iterations/queries at each time instant, both methods can provide a good tradeoff between decreasing the cumulative regret and limiting the human query frequency. Observe that in Fig. 3(a) for  $r = 0.3$ , the local perturbation scheme performs better than the full perturbation scheme for all values of  $N$ , but this advantage becomes weak as the human motion radius  $r$  increases. This is because as  $r$  grows, the relative change in the human distribution becomes large, and so does the relative change between consecutive optimization problems. In the local perturbation scheme, for each optimization problem, more waypoints of the trajectory returned from the previous time instant have to be modified, due to the larger change in the human distribution when  $r$  increases.

### D. Limitations and Future Research Directions

We have conducted a large number of numerical simulations and have tested the proposed framework in challenging planning problems. One limitation of this framework, in its current form, is that it does not perform well in situations where the target location of the robot is very close to obstacles. This is due to the potential functions used for obstacle avoidance in the MPC motion planner that repel the robot trajectory away from the obstacles. For example, our method has not been able to solve planning problems where the target is located in narrow passages formed by obstacles.

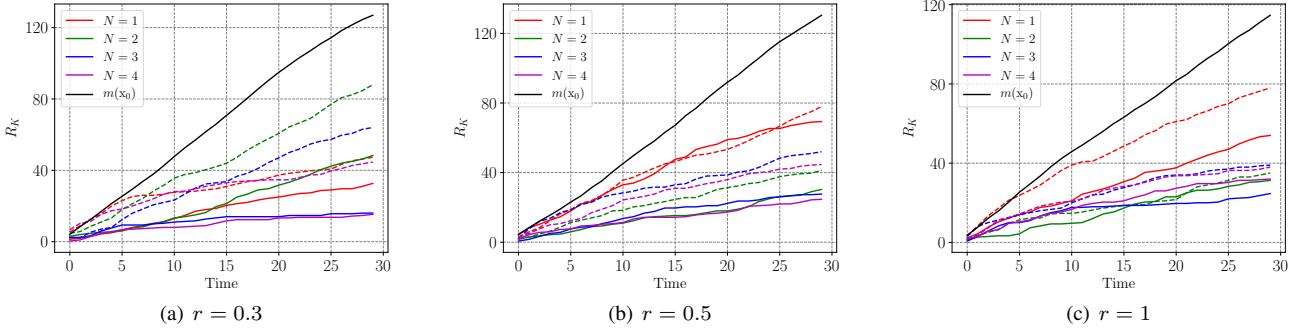


Figure 3. Cumulative regret with respect to  $N$  for different human motion radii  $r$ . The  $x$ -axis corresponds to the index  $i$  of time instant  $T_i$ . The black solid line is the baseline. The colored dashed and solid lines represent the cumulative regret achieved by the full and local perturbation schemes, respectively.

Additionally, the proposed framework in its present form can not easily handle complaints, e.g., about robots being too far from humans; in our simulations complaints are about robots being too close to humans. There are a couple of reasons for this. First, the current local perturbation scheme is not applicable since humans are not able to specify a trajectory segment that affects them if the trajectory is far from them. Second, the full perturbation scheme is also challenging to apply because it is possible that trajectories at consecutive iterations are first close and then far from humans so that it is difficult to estimate zeroth order gradients. The reason for this behavior is the stochastic nature of the full perturbation method. We are currently exploring extensions of our proposed socially-aware planning framework to more complex environments and diverse reasons for human complaints.

## VI. CONCLUSION

In this paper, we considered the problem of designing collision-free, dynamically feasible, and socially-aware trajectories for robots operating in environments populated by humans. We assumed that humans can provide bandit feedback indicating a complaint or no complaint on the part of the robot trajectory that caused them discomfort, and do not reveal any contextual information about their locations or the reasons for their complaints. Assuming that humans can move in the obstacle-free space and, as a result, human utility can change, we formulated this planning problem as an online optimization problem that minimizes the social value of the time-varying robot trajectory, defined by the total number of incurred human complaints. As the human utility is unknown, we employed zeroth order optimization methods to solve this problem, which we combined with off-the-shelf motion planners to satisfy the dynamic feasibility and collision-free properties of the resulting trajectories. To the best of our knowledge, this is the first approach to socially-aware robot planning that is not restricted to avoiding collisions with humans but, instead, focuses on increasing the social value of the robot trajectories. Com-

pared to data-hungry preference learning methods that can be used to solve similar but stationary socially-aware planning problems, our framework is more practical since it only requires small number of iterations and, therefore, queries for human feedback, to return socially-aware robot paths.

## REFERENCES

- [1] T. Kruse, A. K. Pandey, R. Alami, and A. Kirsch, “Human-aware robot navigation: A survey,” *Robotics and Autonomous Systems*, vol. 61, no. 12, pp. 1726–1743, 2013.
- [2] Y. F. Chen, M. Everett, M. Liu, and J. P. How, “Socially aware motion planning with deep reinforcement learning,” in *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2017, pp. 1343–1350.
- [3] G. Ferrer, A. Garrell, and A. Sanfeliu, “Robot companion: A social-force based approach with human awareness-navigation in crowded environments,” in *2013 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2013, pp. 1688–1694.
- [4] P. Trautman, J. Ma, R. M. Murray, and A. Krause, “Robot navigation in dense human crowds: Statistical models and experimental studies of human–robot cooperation,” *The International Journal of Robotics Research*, vol. 34, no. 3, pp. 335–356, 2015.
- [5] M. Kuderer, H. Kretzschmar, C. Sprunk, and W. Burgard, “Feature-based prediction of trajectories for socially compliant navigation,” in *Robotics: science and systems*, 2012.
- [6] H. Kretzschmar, M. Spies, C. Sprunk, and W. Burgard, “Socially compliant mobile robot navigation via inverse reinforcement learning,” *The International Journal of Robotics Research*, vol. 35, no. 11, pp. 1289–1307, 2016.
- [7] A. Vemula, K. Muelling, and J. Oh, “Modeling cooperative navigation in dense human crowds,” in *2017 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2017, pp. 1685–1692.
- [8] A. Alahi, K. Goel, V. Ramanathan, A. Robicquet, L. Fei-Fei, and S. Savarese, “Social lstm: Human trajectory prediction in crowded spaces,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 961–971.
- [9] N. Ratliff, M. Zucker, J. A. Bagnell, and S. Srinivasa, “Chomp: Gradient optimization techniques for efficient motion planning,” 2009.

- [10] J. Schulman, J. Ho, A. X. Lee, I. Awwal, H. Bradlow, and P. Abbeel, “Finding locally optimal, collision-free trajectories with sequential convex optimization,” in *Robotics: science and systems*, vol. 9, no. 1. Citeseer, 2013, pp. 1–10.
- [11] M. Toussaint, “A tutorial on newton methods for constrained trajectory optimization and relations to slam, gaussian process smoothing, optimal control, and probabilistic inference,” in *Geometric and numerical foundations of movements*. Springer, 2017, pp. 361–392.
- [12] N. Ratliff, M. Toussaint, and S. Schaal, “Understanding the geometry of workspace obstacles in motion optimization,” in *2015 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2015, pp. 4202–4209.
- [13] D. J. Webb and J. Van Den Berg, “Kinodynamic rrt\*: Asymptotically optimal motion planning for robots with linear dynamics,” in *2013 IEEE International Conference on Robotics and Automation*. IEEE, 2013, pp. 5054–5061.
- [14] J.-J. Kim and J.-J. Lee, “Trajectory optimization with particle swarm optimization for manipulator motion planning,” *IEEE Transactions on Industrial Informatics*, vol. 11, no. 3, pp. 620–631, 2015.
- [15] M. Kalakrishnan, S. Chitta, E. Theodorou, P. Pastor, and S. Schaal, “Stomp: Stochastic trajectory optimization for motion planning,” in *2011 IEEE international conference on robotics and automation*. IEEE, 2011, pp. 4569–4574.
- [16] S. G. McGill, S.-J. Yi, and D. D. Lee, “Low dimensional human preference tracking for motion optimization,” in *2016 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2016, pp. 2867–2872.
- [17] A. Jain, S. Sharma, T. Joachims, and A. Saxena, “Learning preferences for manipulation tasks from online coercive feedback,” *The International Journal of Robotics Research*, vol. 34, no. 10, pp. 1296–1313, 2015.
- [18] A. Menon, P. Kacker, and S. Chitta, “Towards a data-driven approach to human preferences in motion planning,” in *2015 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2015, pp. 920–927.
- [19] C. Daniel, O. Kroemer, M. Viering, J. Metz, and J. Peters, “Active reward learning with a novel acquisition function,” *Autonomous Robots*, vol. 39, no. 3, pp. 389–405, 2015.
- [20] C. Wirth, R. Akroud, G. Neumann, and J. Fürnkranz, “A survey of preference-based reinforcement learning methods,” *The Journal of Machine Learning Research*, vol. 18, no. 1, pp. 4945–4990, 2017.
- [21] A. Wilson, A. Fern, and P. Tadepalli, “A bayesian approach for policy learning from trajectory preference queries,” in *Advances in neural information processing systems*, 2012, pp. 1133–1141.
- [22] R. Busa-Fekete, B. Szörényi, P. Weng, W. Cheng, and E. Hüllermeier, “Preference-based reinforcement learning: evolutionary direct policy search using a preference-based racing algorithm,” *Machine Learning*, vol. 97, no. 3, pp. 327–351, 2014.
- [23] J. Fürnkranz, E. Hüllermeier, W. Cheng, and S.-H. Park, “Preference-based reinforcement learning: a formal framework and a policy iteration algorithm,” *Machine learning*, vol. 89, no. 1-2, pp. 123–156, 2012.
- [24] P. F. Christiano, J. Leike, T. Brown, M. Martic, S. Legg, and D. Amodei, “Deep reinforcement learning from human preferences,” in *Advances in Neural Information Processing Systems*, 2017, pp. 4299–4307.
- [25] D. Sadigh, A. D. Dragan, S. Sastry, and S. A. Seshia, “Active preference-based learning of reward functions.” in *Robotics: Science and Systems*, 2017.
- [26] R. Pinsler, R. Akroud, T. Osa, J. Peters, and G. Neumann, “Sample and feedback efficient hierarchical reinforcement learning from human preferences,” in *2018 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2018, pp. 596–601.
- [27] Y. Nesterov and V. Spokoiny, “Random gradient-free minimization of convex functions,” *Foundations of Computational Mathematics*, vol. 17, no. 2, pp. 527–566, 2017.
- [28] A. V. Gasnikov, E. A. Krymova, A. A. Lagunovskaya, I. N. Usmanova, and F. A. Fedorenko, “Stochastic online optimization. single-point and multi-point non-linear multi-armed bandits. convex and strongly-convex case,” *Automation and remote control*, vol. 78, no. 2, pp. 224–234, 2017.
- [29] S. Shahrampour and A. Jadbabaie, “Distributed online optimization in dynamic environments using mirror descent,” *IEEE Transactions on Automatic Control*, vol. 63, no. 3, pp. 714–725, 2017.
- [30] A. Mokhtari, S. Shahrampour, A. Jadbabaie, and A. Ribeiro, “Online optimization in dynamic environments: Improved regret rates for strongly convex problems,” in *2016 IEEE 55th Conference on Decision and Control (CDC)*. IEEE, 2016, pp. 7195–7201.
- [31] J. B. Rawlings, “Tutorial overview of model predictive control,” *IEEE control systems magazine*, vol. 20, no. 3, pp. 38–52, 2000.
- [32] L. Grüne and J. Pannek, “Nonlinear model predictive control,” in *Nonlinear Model Predictive Control*. Springer, 2017, pp. 45–69.
- [33] F. Allgöwer and A. Zheng, *Nonlinear model predictive control*. Birkhäuser, 2012, vol. 26.
- [34] L. Lapierre, R. Zapata, and P. Lepinay, “Combined path-following and obstacle avoidance control of a wheeled robot,” *The International Journal of Robotics Research*, vol. 26, no. 4, pp. 361–375, 2007.
- [35] Y. Yoon, J. Shin, H. J. Kim, Y. Park, and S. Sastry, “Model-predictive active steering and obstacle avoidance for autonomous ground vehicles,” *Control Engineering Practice*, vol. 17, no. 7, pp. 741–750, 2009.
- [36] Y. Gao, T. Lin, F. Borrelli, E. Tseng, and D. Hrovat, “Predictive control of autonomous ground vehicles with obstacle avoidance on slippery roads,” in *ASME 2010 dynamic systems and control conference*. American Society of Mechanical Engineers Digital Collection, 2011, pp. 265–272.
- [37] L. Zhang, T. Yang, J. Yi, J. Rong, and Z.-H. Zhou, “Improved dynamic regret for non-degenerate functions,” in *Advances in Neural Information Processing Systems*, 2017, pp. 732–741.
- [38] S. Karaman and E. Frazzoli, “Sampling-based algorithms for optimal motion planning,” *The international journal of robotics research*, vol. 30, no. 7, pp. 846–894, 2011.
- [39] C. J. Ostafew, A. P. Schoellig, and T. D. Barfoot, “Robust constrained learning-based nmpc enabling reliable mobile robot path tracking,” *The International Journal of Robotics Research*, vol. 35, no. 13, pp. 1547–1563, 2016.