

Risk-Based Human-Aware Multi-Robot Coordination in Dynamic Environments Shared with Humans

Zeynab Talebpour and Alcherio Martinoli

Abstract—In this paper, we propose a risk-based coordination method for the Multi-Robot Task Allocation (MRTA) problem in human-populated environments. We introduce risk-based bids that incorporate human trajectory prediction uncertainties and furthermore, social costs in their formulation. We demonstrate the effectiveness of including a predictive component in the risk formulation despite the lack of accurate position estimation for humans through an extensive suite of experiments. This is done by means of testing different levels of prediction error for known human trajectories and in a separate approach, using a Kalman filter for human trajectory estimation. Furthermore, we propose different risk formulations and evaluate their performance in a high-fidelity simulator. Additionally, a comparative study targeting human-agnostic planning at both navigation and planning levels, human-aware navigation and planning based on deterministic costs, and risk-based human-aware planning with no individual human-aware navigation has been conducted. Results confirm that risk-based bids lead to more socially acceptable team plans that reduce the need for the lower level individual human-aware navigation to be activated. Risk-based plans that account for social costs, prevent difficult social situations that can lead to less effective human-aware navigation, such as traversing narrow passages occupied by humans.

I. INTRODUCTION

With the increasing presence of technology and robots in social environments, research in the area of socially-aware navigation has received much attention in the recent years. Robots in social environments either directly interact with people or have to operate in spaces shared with humans. Therefore, it is essential that robots gain the acceptance of people by ensuring their safety and comfort [1] and demonstrating legible and predictable [2] actions.

Despite the numerous applications of Multi-Robot Systems (MRS) in social environments, such as personal assistants at homes, robot tutors at schools, service robots at hospital and nursing homes, research in the human-aware navigation area focuses mainly on single robots and the problem of cooperative human-aware navigation for MRS -an interesting problem for both multi-robot and human-aware navigation research- is largely unexplored.

In this paper, we focus on a particular class of MRS coordination mechanisms commonly known as Multi-Robot Task Allocation (MRTA) [3], [4] in social environments using a market-based approach [5]. In such environments the number of robots are often limited and the number of tasks are usually moderate. The main difficulty for MRTA in such highly dynamic and noisy environments, is that plans are likely to change or to be rendered invalid, particularly if the robots are planning for long periods of time. Additionally,

Zeynab Talebpour and Alcherio Martinoli are with the Distributed Intelligent Systems and Algorithms Laboratory (DISAL), School of Architecture, Civil and Environmental Engineering, École Polytechnique Fédérale de Lausanne (EPFL), 1015 Lausanne, Switzerland, `firstname.lastname@epfl.ch`




Fig. 1: Two robots bidding on a task that will have a changing cost over time. The human is initially static. He then decides to move towards Robot 1. This snapshot shows the initial position of the robots and the human for case study 1 in Webots.

robots are required to perform in a socially acceptable manner in terms of navigation and interaction with people and other team members. This adds additional constraints to the planning problem.

In social environments, when robots take decisions based on the current available information only, they are assuming that their decision remains valid in its period of execution. However, human actions such as walking, starting or ending an interaction, can largely modify the social costs while a robot is proceeding towards an already allocated task. In such cases, the initial bid estimation for the task allocation is no longer a true representative of the real cost. As a result, the performance and efficiency of team plans can degrade in more complex and dynamic environments. In this paper, we study MRTA in dynamic social environments with costs that have a stochastic nature due to the changing behavior of people.

To illustrate the problem, we will use the example in Fig. 1, where two robots are coordinating to find the best team plan for taking Task 1. In a social MRTA approach that only considers the current available information when bidding on a task and given a local balance function that scores tasks inversely proportional to their distance to the robot, Robot 1 will take Task 1. This is because the distance that Robot 1 has to travel (d_1) is smaller than the distance traveled by the other robot (d_2). Additionally, Robot 2 will have to travel an even larger distance (d_3) to avoid the social costs associated to the personal space of the human. However, in reality the human is moving towards Robot 1. This means there will be no social costs associated to the path of Robot 2. Additionally, Robot 1 will have to modify its path in order to respect the personal space of the human and will have to travel a larger distance than planned. If robots consider the future positions of the human while estimating their bids, they will know that Robot 2 is a better candidate for taking Task 1 in this situation.

Accurate estimation of future positions of humans in uncontrolled environments is not possible in reality. Nonetheless, prediction of human motion despite being error prone, can still provide valuable information about the changes that are likely to occur. Taking a decision based on uncertain information can be seen as taking a risk. In the context of MRTA, a risk

measure that adopts predictions and captures their errors as uncertainties can be a useful extension to the deterministic estimation of costs. Furthermore, accounting for the added social costs corresponding to risky situations, can help the robots to take more informed and socially-aware decisions.

The contributions of this work include, proposing a concept of risk-based human-aware bids that account for changing costs in human-populated environments, evaluating the effect of human trajectory prediction error on risk-based bid estimation and team performance, and proposing different risk formulations to account for the prediction error, risk estimation accuracy, and social costs. To the best of our knowledge, MRTA with stochastic human-aware costs in social environments has not been investigated in the literature.

II. RELATED WORK

MRTA algorithms vary in design and application [6], [7], but their common objective is to find a mapping between robots in a team and a set of “tasks” that must be accomplished in order for the team “goal” to be completed. Among multiple approaches proposed for MRTA, we are mainly interested in distributed approaches that can be executed by a team of robots without the explicit need for a centralized entity outside the team with perfect knowledge of the environment [8].

Market-based multi-robot coordination [5], [9], [10] is an example of such a MRTA approach. In these systems, robots act as agents trying to maximize their individual profits. Every time a task is auctioned, robots must pay a price to obtain it. Once the task is completed, a payment is done to the robot which won the auction. The underlying assumption is that with every robot trying to maximize its individual profit, the overall team coordination and efficiency will be improved.

As stated in [6], despite the importance of uncertainty in real robotic problems and the potential of stochastic planning for producing sound and robust allocation policies, most MRTA approaches assume a deterministic MRTA model and deal with uncertainty only at execution time by replanning during task execution.

In stochastic allocation literature, it is assumed that a model of uncertainty, for instance a probability distribution of robot travel time, task arrival, etc., is available. Such MRTA problems are commonly modeled as Markov Decision Processes (MDPs) [11], or as pure or mixed stochastic integer programs [12]. In a different approach, approximation of the parametric uncertainties captured by the underlying system model has been investigated in [13] by means of active learning.

Despite the research in this area, it is not clear what is the better approach for facing uncertainty and the challenges of MRTA still remain open. In [6] such dilemma is summarized as follows: “Is it more beneficial to build a complex model that incorporates uncertainty, or is it enough to build less well-informed plans and replan as often as needed to quickly react to unexpected events?”

Planning under uncertainty can also be addressed using Partially Observable Markov Decision Processes (POMDPs). However, this approach is faced with a scalability problem when considering teams of robots. Auctioning of independent local POMDP-based controllers is proposed by [14] to alleviate this problem. Nonetheless, for real uncontrolled environments, the information space get too large to be tractable.

Very little work has been done in MRTA for social environments. In [15], a general MRS architecture for person search for a team of assistive robots in a retirement home is proposed. Therein, MRTA is considered as a constraint optimization problem and a centralized planning approach based on constraint programming [16] is used to solve it. The focus of this work is on the capabilities of the proposed architecture and social aspects have not been considered. In [17], a real-world multi-robot coordination problem for human guidance, requiring stochastic transitions is shown to be successfully implemented in a centralized fashion, at a scale of five to ten robots. The human behavior is simplified and a discrete state space representing key locations on a map is used for the MDP. Similarly, no social costs have been considered in this study.

In our previous work, a study of MRTA in dynamic and noisy environments [18] for spatial task allocation confirmed the effectiveness of our Hoplitess-based method [19] as a first step towards deployment of MRTA methods in social environments. In [20], we proposed a human-aware coordination method for MRTA that accounted for the distance overhead of human-aware paths in the local balance function of the robots and allowed for instant collaboration in socially blocking situations. However, robots could only see a snapshot of the environment and costs were assumed to be deterministic.

MRTA in social environments should ensure social acceptance as well as achieving good performances in terms of global team objectives such as the traveled distance, mission time, etc. Respecting personal spaces, O-spaces and P-spaces [21] are the common social behaviors considered in the literature of human-aware navigation based on the concept of “proxemics” [22]. In this work, proxemics-based social costs encoded as costmaps similar to [23] along with a Fast Marching Method (FMM)-based path planner constitute the human-aware navigation components of the robots. FMM has been proven to be successful in real domestic spaces with high complexity [24]. There have previously been a number of research papers addressing social path planning using FMM [25].

Unlike most stochastic MRTA approaches, an uncertainty model for uncontrolled social environments is not available unless strong assumptions are made or a data-driven approach, targeting a specific environment, is taken. Additionally, the scale and complexity of the problem is too large for applying POMDP-based solutions. Moreover, each encounter of the robots with humans matters and improving the average performance is not the best strategy to gain social acceptance for the robots. As a result, we opt for an approach that uses risk as a heuristic for estimating stochastic costs.

Risk-based navigation (Risk-RRT [21]) has been adopted in human-aware navigation for single robots. Therein, risk, or the probability of collision with objects or intrusion in socially costly areas, along any candidate trajectory is taken into account for selecting an appropriate human-aware path. Inspired by this idea, we propose a number of risk formulations for estimation of stochastic costs for social MRTA in the next section.

III. PROBLEM DESCRIPTION AND APPROACH

In this section we define the MRTA problem and explain the underlying market-based framework for robot coordination.

Next we propose the risk-based bid formulations.

A. Problem Description

Consider a team of n_r robots $\{r_j, j = 1, \dots, n_r\}$ and a set of n_t tasks $\{t_i, i = 1, \dots, n_t\}$, where each task is a location that a robot has to visit. The team of robots should decide how to efficiently subdivide these tasks using a local criterion in order to optimize a global performance. This global performance can be a function of time, distance traveled, etc. Tasks can be identified locally by the robots through on-board perception or can be broadcasted to all robots by an external source. Many applications such as patrolling, attending service requests, etc., can benefit from this functionality.

For succeeding in a social environment, robots are required to ensure that social constraints, namely personal and interaction spaces of humans, are not violated.

B. Social Multi-Robot Coordination

By means of a Hoplitest-based coordination [19] scheme comprised of two main concurrent types of coordination, *passive coordination* and *active coordination*, the robots, which are self-interested agents in pursuit of individual profit, can evaluate each available task and decide whether to take it or sell/buy it to/from another robot in a distributed manner.

In **passive coordination**, each robot chooses its most profitable plan and broadcasts it to other teammates without any attempt to modify their plans. This information is then used by other robots to reevaluate the expected profitability of their current plans, update and broadcast the changes. However, sometimes a robot's best plan can only be marginally profitable while a genuine team plan could result in a higher profit. This suggests that modifying the plans of the robot's teammates would be an interesting option to pursue. Concretely, this implies that the requesting robot asks its collaborators for compensation price quotes and persuades them to engage in a cooperative action. This process is ruled by a market-based approach and constitutes the **active coordination**. For details about the Hoplitest-based MRTA refer to [18] and [20].

The decision of switching to the active coordination mode is based on the evaluation of a **local balance function**. This function is strongly problem-dependent and can contribute to reaching the globally optimal solution at the team level if chosen correctly. Additionally, a problem-dependent **global balance function** is also required for team-level evaluations.

The local balance functions constitute what is known as a bid in the context of the market. While many formulations can be used for encoding the local balance function of robots, operating in social environments calls for including a social factor in bids. Thus, human-aware bids and social time-outs are used to encode the social information into the coordination framework [20]. Social time-outs are used for requesting collaboration when socially blocking situations occur.

C. Deterministic Social Bids

Robots should find the most appropriate assignment for optimizing the global balance function based on a local balance function. The local balance function of robot r_j , for each task t_i belonging to a plan P at time k is defined in the following. This function is inversely proportional to the length of the path planned by the FMM to the desired

task. We note that for social task planning, the path planned by the FMM already takes into account the social costmaps representing the human-centric Gaussian cost functions [26].

$$B_{j,P}(k) = \sum_{t_i \in P(k)} (R_{t_i}(k_{r,t_i}) - D_{social}(l_{t_{i-1}}, l_{t_i})) \quad (1)$$

R_{t_i} is the revenue of task t_i , l_{t_i} is the position of t_i and $l_{t_{i-1}}$ is the position of the task that appears before t_i in P . k_{r,t_i} is the time that r_j is expected to reach t_i and D_{social} is the distance traveled based on the human-aware path. This balance function includes a revenue R_{t_i} that is decreasing with time as shown below.

$$R_{t_i}(k_{r,t_i}) = \max(0, R_{max}(1 - \frac{k_{r,t_i} - k_{a,t_i}}{T})) \quad (2)$$

where k_{r,t_i} is the time in which t_i is reached, R_{max} is the maximum revenue for the task, k_{a,t_i} is the allocation time of t_i and T is the time after which the positive revenue becomes zero. This utility function is added to reinforce reaching the tasks as early as possible. In real noisy social environments, only an estimate of the time to reach a task (k_{r,t_i}) can be computed and due to the changing behavior of people, estimates of k_{r,t_i} and the distance traveled by the social robot (D_{social}) can be far from accurate for some of the robots and acceptable for the ones who are not faced with humans. If a static person starts moving and clears the robot's path, the initial bid on the task has been an over-estimation. On the contrary, if a moving human suddenly occupies parts of a robot's path or decides to interact with the robot, the original bid has been an under-estimation. In other words, costs associated with tasks are uncertain and can vary over time.

D. Risk-based Bids

Accurate prediction of the future for humans is not possible because of the uncertainties inherent to uncontrolled social environments. Therefore, we propose an abstraction that can extract higher level information from perceptual data, by introducing risks. "Risk" is defined as the probability of occupation of an area with social costs by the robot.

In general, the local balance function B for a robot r_j and a given plan P at time k is defined as follows:

$$\begin{aligned} B_{r_j,P}(k) &= R_{r_j,P}(k) - C'_{r_j,P}(k) \\ C'_{r_j,P}(k) &= C_{r_j,P}(k) + G_{r_j,P}(k) \end{aligned} \quad (3)$$

where R is a generic revenue function, C a generic cost function and G the penalty for constraint violations. In a human-populated environment, the presence of people can lead to an increase in the cost for a plan compared to the estimated cost of that plan in an empty deterministic and noise-free environment. Similar to the way humans assess situations by means of evaluating the risks against the benefits when taking a decision, the robot will compute a risk-based cost C'' that will be aggregated with a revenue R for a given plan P .

Another term Q is introduced to captures this stochastic cost which is proportional to the risk γ associated with the plan. f_m is a user-defined function used to aggregate the risk with revenue, cost and penalty terms, based on the risk formulation method m .

$$\begin{aligned}
B_{r_j, P(k)} &= R_{r_j, P(k)} - C''_{r_j, P(k)} \\
C''_{r_j, P(k)} &= C'_{r_j, P(k)} + Q_{r_j, P(k)} \\
Q_{r_j, P(k)} &= \sum_{t_i \in P(k)} f_m(\gamma_{r_j, t_i, k})
\end{aligned} \tag{4}$$

When evaluating B_{r_j} , the risk of being subjected to additional costs due to human actions for each task t_i in P throughout the mission must be determined. This risk is computed on the basis of the distance D between the robot and any human present in a predefined vicinity of the robot at any point of time. Let's assume that at time k , t_i will be started in k_1 seconds and will be reached in k_2 seconds. The risk associated with this part of the plan is defined as follows:

$$\gamma(r_j, t_i, k) = \int_{k+k_1}^{k+k_2} \sum_{h \in H} g_m(D(l_{r_j, k'}, l_{h, k'})) dk' \tag{5}$$

Here, g_m is a function that is inversely proportional to D , i.e., the distance measure between the robot position l_{r_j} and the position of the human l_h . Since risk is defined as the probability of occupation of an area with social costs by the robot, g is chosen to normalize the distance in a predefined vicinity of the human with radius ϵ_R :

$$g_m = \frac{\max\{0, \epsilon_R - D_m\}}{\epsilon_R + D_m} \tag{6}$$

The key information for computing risk is the future positions of people. This calls for a human trajectory predictor to make the required information available. Throughout this paper, we assume that the output of the human trajectory predictor is a Gaussian distribution. $M_{h, k}$ is the mean and $\Sigma_{h, k}$ is the covariance matrix of the distribution. For time k and human h , we will have a predictor that provides us with the following information:

$$l_{h, k} = \mathcal{N}(M_{h, k}, \Sigma_{h, k}) \tag{7}$$

Robot position on the other hand, can be given by the path planner. For any given goal (l_{t_i}), we obtain a sequence of way-points $w \in W_{r_j, t_i} = \{w_1, w_2, \dots, w_N\}$ that constitute the robot path, with w_0 being the initial position of the r_j . An estimate of the corresponding reaching time for each w is also known $k_w \in \{k_{w_1}, k_{w_2}, \dots, k_{w_N}\}$. The robot trajectory is then discretized and k is ignored as it is implicitly accounted for through the use of k_w . The new risk formulation can be written as:

$$\gamma(r_j, t_i) = \sum_{w \in W_{r_j, t_i}} \sum_{h \in H} g_m(D(w, l_{h, k_w})) \tag{8}$$

H represents the set of all humans perceived by the robot. Depending on the risk modeling approach, different distance and risk formulations can be chosen. We are interested to know how this choice can affect the performance of the team plan in terms of both MRTA and social metrics. Therefore, we propose the following risk formulations. Details of our proposed generic bid estimation algorithm can be found in Algorithm 1.

1) *Euclidean Distance*: By choosing this metric, using a simple formulation that requires only the end points without the uncertainty associated with them we compute the straight-line distance between two points in Euclidean space, i.e., way-point and human positions. The expected value of the human prediction distribution will be used as the position of the human.

$$\gamma(r_j, t_i) = \sum_{w \in W_{r_j, t_i}} \sum_{h \in H} g_m(D_E(w, \mathbb{E}(l_{h, k_w}))) \tag{9}$$

2) *Mahalanobis Distance*: Since human trajectory predictions provide a Gaussian distribution with a covariance matrix, we do know about the uncertainty associated with every prediction. Thus, a more accurate distance measure can be extracted if this uncertainty is taken into account by means of Mahalanobis distance.

The Mahalanobis distance between robot r_j and human h can be written as follows. Note that $l_{r_j} = w$.

$$D_M(l_{r_j}, l_{h, k}) = \sqrt{(l_{r_j} - M_{h, k_w})^\top \Sigma_{h, k_w}^{-1} (l_{r_j} - M_{h, k_w})} \tag{10}$$

$$\gamma(r_j, t_i) = \sum_{w \in W_{r_j, t_i}} \sum_{h \in H} g_m(D_M(w, l_{h, k_w})) \tag{11}$$

3) *Integrated Distances*: In sections III-D.1 and III-D.2, only way-points from the planner have been considered for risk computation. However, if the granularity of robot path is not fine enough, there are cases where important events can be missed; for instance, robot and human collision or social zone intrusion can happen in between two way-points without the two ends of the trajectory segments being affected by it. Therefore, a piece-wise linear breakdown of the segment between way-points can be considered with a predefined resolution. The formulation of g_m for D_E and D_M for integrated distances can be found in the following:

$$g_m(D_E(w, \mathbb{E}(l_{h, k_w}))) = \max_{w_p \in PWL(w_{-1}, w)} g_m(D_E(w_p, \mathbb{E}(l_{h, k_{w_p}}))) \tag{12}$$

$PWL(w_{-1}, w)$ represents the piece-wise linear breakdown of the segment between w and its previous waypoint w_{-1} . Instead of max operator, summation can also be used to account for the accumulated social costs. Similarly, for D_M integrated risk will make use of the following distance formulation.

$$g_m(D_M(w, l_{h, k_w})) = \max_{w_p \in PWL(w_{-1}, w)} g_m(D_M(w_p, l_{h, k_{w_p}})) \tag{13}$$

4) *Social Cost Incorporation*: Human presence can affect the cost for a task in a number of ways. The distance traveled will increase, risk of interactions and therefore, incomplete missions are introduced and additionally costs concerning human discomfort or inconvenience can be incurred. By assigning costs on the basis of social costmaps to the risk formulation, social factors are further reinforced.

Expectation-based social costmaps are used to incorporate the uncertainty in the human positions reported by the human trajectory predictor. Consider a person at (x_p, y_p) , the deterministic costmap at (x, y) is:

$$S(x, y; x_p, y_p) = N(x - x_p, y - y_p) \tag{14}$$

N is the 2D Gaussian modeling the standard social costmaps. The probabilistic costmap is given by the expected value of the social cost S , given the probability distribution of the human's position $p(x_p, y_p)$ which in this case belongs to a Normal distribution for any point in its vicinity, e.g. $w = [x, y]^\top$ is:

$$\mathbb{S}(w) = \mathbb{E}(p(x_p, y_p)[S(x, y; x_p, y_p)]) = \int \int N(x - x_p, y - y_p) p(x_p, y_p) dx_p dy_p \tag{15}$$

This is in fact the convolution of the Normal function modeling the human position l_{h,k_w} and a Gaussian social costmap model S centered on the mean of l_{h,k_w} which is $[x_p, y_p]^T$. Using this approach, we compute an *expected costmap* incorporating all the uncertainty in the environment. Based on this social costmap we now can have a risk value that incorporates social costs directly. The risk formulation for endpoint and integrated social costs can be written as the following respectively:

$$\gamma(r_j, t_i) = \sum_{w \in W_{r_j, t_i}} \sum_{h \in H} S(w) \quad (16)$$

$$\gamma(r_j, t_i) = \sum_{w \in W_{r_j, t_i}} \sum_{h \in H} \max_{w_p \in PWL(w_{-1}, w)} S(w_p) \quad (17)$$

The potential additional social costs only occur in the vicinity of humans. This means for scenarios with many tasks and many robots, only a subset of robots and tasks which are subjected to social costs are making the difference with non risk-based approaches. The building block of such scenarios, is bidding on one task affected by one human. If robots could improve their estimates of the cost for such a task, they will subsequently improve their team level performance since more grounded decisions will be taken. Additionally, by means of an aggregated formulation of risk that accounts for all the humans perceived by the robot, every human is considered when computing the bid and areas containing more humans will be associated a higher risk.

There are two main components to MRTA performing well in dynamic social environments: 1) improved bid estimation, 2) adaptive replanning. In this paper, we focus on the former since we believe providing a detailed evaluation of different methods for risk-based bid computation is the first essential step for understanding how to approach stochastic social costs for acquiring better team plans. Moreover, reliable risk estimation is the basis for devising adaptive replanning strategies that can accommodate the high dynamics of social environments. Currently, re-planning is done when a task is accomplished or for verifying the validity of a stored plan when a robot is on its way towards a task (refer to [18] for detail).

Algorithm 1 Risk-based bid estimation of robot r_j for task t_i with humans H

```

1: procedure BIDESTIMATION( $l_{r_j}, l_{t_i}, H$ )
2:    $\triangleright$  Compute the path to  $t_i$  assuming an empty map
3:    $(W_{r_j, t_i}, K_{r_j, t_i}, pathCost_{r_j, t_i}) \leftarrow$  PATHPLANNING( $l_{r_j}, l_{t_i}$ )
4:    $bid_{r_j, t_i} \leftarrow pathCost_{r_j, t_i}$ 
5:    $w_{-1} \leftarrow l_{r_j}$ 
6:   for  $w \in W_{r_j, t_i}$  do
7:     for  $h \in H$  do
8:        $\triangleright$  Find human position at time  $k_w$ 
9:        $l_{h, k_w} \leftarrow$  PREDICTHUMANPOSITION( $l_{h, k_{w-1}}, k_w$ )
10:       $\triangleright$  Estimate the risk of  $h$  for  $r_j$  between  $w_{-1}$  and  $w$ 
11:       $\gamma_{r_j, h, w} \leftarrow$  ESTIMATERISK( $w_{-1}, w, l_{h, k_{w-1}}, l_{h, k_w}$ )
12:       $\triangleright$  User-defined function for aggregating risks with the bid
13:       $bid_{r_j, t_i} \leftarrow$  AGGERATERISK( $\gamma_{r_j, h, w}, bid_{r_j, t_i}$ )
14:    $w_{-1} \leftarrow w$ 
return  $bid_{r_j, t_i}$ 

```

IV. METRICS AND NAVIGATION

In this section we will describe the evaluation metrics, navigation method, the robots used for our experiments, and the simulation tool.

A. Metrics

For the global balance function concerning MRTA, the total traveled distance (M_1) and the mission time (M_2) are reported for all experiments. For evaluating the performance of the MRS in terms of social-awareness, the maximum accumulated social cost (M_3), the maximum time steps spent in areas associated with social costs (M_4) and the minimum distance to any human throughout the experiment (M_5) are reported among all robots.

B. Navigation

The navigation system is that of the MONarCH project [27], detailed in [24]. As input, it uses the pose estimates provided by a standard Adaptive Monte Carlo Localization (AMCL) self-localization system, given odometry, laser range finder readings, and a static map. The navigation system is based on FMM for motion planning, together with a Dynamic Window Approach (DWA) algorithm for guidance and obstacle avoidance. FMM and DWA run asynchronously. FMM is activated when a new goal position is given, and DWA is running in a closed loop with a fixed rate of 20 Hz in our experiments.

C. Robot

The robot we used in this work is called the MBot [28] (see Fig. 1) and has been developed within the FP7 European project MONarCH¹. It is an omni-directional drive robot with an approximately round footprint of 0.65 m in diameter and a height of 0.98 m. It is endowed with two laser range finders, on both the front and the back for providing 360° coverage.

D. Simulations

The use of high-fidelity simulators such as Webots [29] is fundamental, especially when considering multi-robot systems. When evaluating MRS in human-populated environments, conducting real experiments with varying features and complexity levels becomes exceeding difficult. We have developed models of the environments (see Fig. 1 and Fig. 2) for our experiments that enable simulations with similar environmental richness to the real world experiments, as shown in [18].

V. EXPERIMENTS

Two case studies of increasing complexity will be described in this section. In section V-A, we will investigate the performance of different risk formulations as well as the effect of human trajectory prediction error on the method performance. In section V-B, a comparative study targeting 1) human-agnostic navigation and planning, 2) human-aware navigation without considering humans in the planning phase, 3) human-aware navigation and planning based on deterministic costs, and 4) human-aware planning based on stochastic costs without any individual human-aware navigation will be explained.

We note that in each case study, human trajectories are the same across all runs with different methods. This choice has been made to ensure that we are comparing the same MRTA problem instance. Each scenario has been repeated for ten simulation runs. Robots are relying on their self-localization for computing the local balance functions and the evaluation metrics ($M_1 - M_5$) have been obtained from ground truth values provided by the simulation.

¹<http://monarch-fp7.eu/>

Scenario	No. of Humans	Risk Method	Noise σ (m)
1A	0	-	-
1B-AG	1	-	-
1B-E0	1	Euclidean	0.0
1B-E0.5	1	Euclidean	0.5
1B-E2	1	Euclidean	2.0
1B-E-KF	1	Euclidean	KF
1B-M0.5	1	Mahalanobis	0.5
1B-M2	1	Mahalanobis	2.0
1B-M-KF	1	Mahalanobis	KF
1B-S0	1	Social costmap	0.0
1B-S0.5	1	Social costmap	0.5
1B-S2	1	Social costmap	2.0
1B-S-KF	1	Social costmap	KF

TABLE I: Details of different scenarios in case study 1. Noise stands for the human trajectory prediction noise.

For computing social metrics, we rely on Gaussian social costmaps. Similar to [23], we have chosen $\sigma_x = 0.255 m$, and $\sigma_y = \sigma_x$ since we don't include orientation in our predictions. The value of social cost for a given position varies between 0 and 100 and the radius of the social costmap is 1 m.

A. Risk Formulation and Trajectory Prediction Analysis

This test case consists of two robots, one task and a dynamic human that causes the cost of the task to change over time (see Fig. 1). The goal of this case study is to investigate the main challenging aspect of changing costs, i.e., impact of the human. This test case serves as the building block for more generalized cases with more robots and more humans. In the presence of multiple people, each robot will perform the same operation to compute an accumulative risk accounting for every perceived human.

For this test case, we have only included the point-based risk methods, since with 0.25 m granularity of the way-points in robot path planning, there is no need for opting for an integrated risk model in our problem.

To study the effect of prediction error on the performance of our method, different levels of noise have been tested for each risk formulation. This is done by means of adding noise to the ground truth trajectory of the human, i.e., sampling a point l_h from a Gaussian distribution $\mathcal{N}(M_{GT}, \Sigma_{Noise})$ with:

$$M_{GT} = [x_{GT}, y_{GT}]^T, \quad \Sigma_{Noise} = \begin{bmatrix} \sigma^2 & 0 \\ 0 & \sigma^2 \end{bmatrix} \quad (18)$$

x_{GT} is the x component and y_{GT} is the y component of the true human position. The output of each prediction is another Gaussian distribution $\mathcal{N}(l_h, \Sigma_{Noise})$. When adopting a Kalman Filter (KF) for human position prediction, a state consisting of $[x, y, v_x, v_y]^T$ is tracked and $[x, y]^T$ is observed. x is the x component of human position in the world frame, y is the y component of human position, and v_x, v_y are the corresponding velocity vector components. A constant velocity dynamics model is assumed for the human motion. Table I lists the details of each scenario in terms of noise-level and risk method.

B. Comparative Evaluation of Different MRTA Strategies

In this case study, a problem consisting of five tasks, four robots and two dynamic humans has been considered. Fig 2 illustrates the initial position of the robots, task placement and human trajectories. The aim of this assessment is to study the performance of the risk-based social task planner in comparison to the other approaches in a more complex setting. This case study demonstrates how team plans can




Fig. 2: Initial position of the robots and humans in test case 2 in Webots. Blue circles depict the location of tasks and smaller circles on human trajectories represent brief pauses.

change if future risks are to be considered and how social costs can be significantly reduced by means of team plans that avoid social risks.

To compare the performance of different MRTA strategies, four series of tests across two scenarios have been conducted. In *Scenario 2A*, there are no humans present, whereas in *Scenario 2B*, two moving humans exist in the environment as explained in the case study description. This latter scenario adopts three different algorithms, *Scenario 2B-AG*, human-agnostic robots, in this scenario humans are not considered in the team plan and are only considered as obstacles in robot navigation. *Scenario 2B-SD*, social deterministic costs and individual human-aware navigation; here humans are considered in task planning but through a social planner that takes decisions based on currently available information only. *Scenario 2B-SR-KF*, risk-based social planning without individual human-aware navigation; in this scenario a KF predictor is used for bid estimation. Among different risk formulations, S-KF was shown to be the most effective (see section VI-A), hence, we selected this method for further evaluation. We chose to decouple the risk-based social planner from individual human-aware navigation to highlight the strengths of this planning approach and assess the contribution of risk-based bids in finding better plans and reducing social costs.

VI. RESULTS AND DISCUSSION

In this section the results of the two case studies explained in the previous section will be discussed.

A. Risk Formulation and Trajectory Prediction Analysis

Consider Fig. 1 again. Upon initial bidding of both robots, Robot 1 is the closest to Task 1 and will therefore be allocated to this task in Scenario 1A and 1B-AG. However, when future risks are taken into account in task planning, in all other scenarios Robot 2 will request a collaboration and take Task 1. Social task planning based on deterministic costs has not been included in this test case since it is clear that with the added distance overhead of social paths (see d_3 in Fig. 1), Robot 2 will not be able to overbid Robot 1. For this test case, plans are not included for brevity and we focus on the extracted metrics instead. Fig. 3 shows the performance of this test case across different scenarios. For Scenario 2A, we expect to have the smallest M_1 and M_2 due to no humans existing in the environment. This is confirmed by Fig. 3a-b. Scenario 1B-AG has similar performance to 1A in terms of $M_1 - M_2$ on average. However, without individual human-aware navigation we observed Robot 1 having a mild or major collision with the human in a few runs. This resulted in pauses and localization errors for the robot and hence the increased variability in M_1 and M_2 compared to Scenario 1A, despite following the same task assignment. Lack of

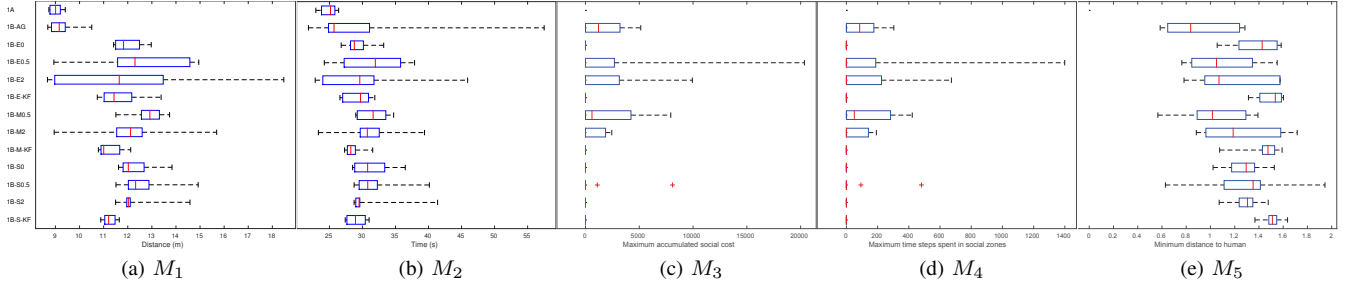


Fig. 3: Performance metrics for the first set of experiments obtained from 10 runs. Note that for metric M_3 to M_5 , no performance for scenario 1A is plotted since no human was present.




Fig. 4: Performance metrics for the second set of experiments obtained from 10 runs. Note that for metric M_3 to M_5 , no performance for scenario 2A is plotted since no human was present.

human-awareness has lead to large social costs for Scenario 1B-AG as seen in Fig. 3c-e.

For risk-based bidding methods, we can observe zero socials cost on average and they maintain an appropriate distance to the human across all tests despite noisy predictions. The only exception here is 1B-M0.5 where Robot 2 was able to persuade Robot 1 later than other scenarios in five of the runs. As a result, Robot 1 stopped at a position that would be sometimes too close to the human trajectory and without activating human-aware navigation or social collision avoidance, Robot 1 ended up being in the human's social zone. Similar behavior was observed for fewer runs in Scenario 1B-E0.5 and Scenario 1B-E2.

All scenarios with zero prediction error, KF predictor, and risks that are based on social costmaps, have resulted in zero social costs. We believe the reason is that social costs reinforce the risk in areas closer to the human, i.e., a larger penalty (based on the social costmap model) is assigned to all areas associated with possible human presence. Despite large prediction errors, particularly for KF, we observe that the correct decision has been taken by the robots. We believe that extracting a direction from predictions is the key for improving the bid estimates here. This is an abstraction compared to accurate position estimation. However, this information can be used to improve team plans in many cases. Based on the results of this case study, we choose the risk formulation that was based on social costmaps with a KF human trajectory predictor for our risk-based bid estimations in the next test case. We note that other scenarios with frequent sudden changes in human trajectories or more complex dynamics can be found where this solution alone will not be sufficient. This is where monitoring risk variations and adaptive replanning can help. We will pursue this further in our future work.

B. Comparative Evaluation of Different MRTA Strategies

Looking at Fig. 2, for a non risk-based task planning strategy, movement of Human 1 and Human 2 causes change to the initial estimated costs of Room 1 and Room 6, respectively. Sample team plans for all four scenarios are shown in Fig. 5. We can see how the team plan changes for Scenario 2B-SR-KF compared to the other scenarios. For a human-free environment where tasks are scored inversely proportional to their distance to the robot and given the proximity of robots to tasks, Room 1 is the best choice for Robot 1 and Room 6 is the best option for Robot 4. This is confirmed by the team plans (see Fig. 5a and 5b) for both Scenario 2A and Scenario 2B-AG. In Scenario 2B-SD humans are taken into account when computing bids by means of including the distance overhead of the social paths. We can see that despite this consideration, the team plan does not change (see Fig. 5c) and instant human-aware decisions are not able to find the more appropriate task assignment. Scenario 2B-SR-KF on the other hand, computes risk-based bids and includes future estimates of human motion in decision making. Thus, a different and less socially intrusive plan is found for the team as depicted in Fig. 5d.

Fig. 4 demonstrates M_1 - M_5 for Scenario 2A to 2B-SR-KF. It can be seen that Scenario 2A has the smallest M_1 - M_2 due to lack of human presence. Slightly larger M_1 - M_2 can be seen for 2B-AG and 2B-SD with large variations in M_2 . This is again caused by the minor or major collisions with humans. We observed that despite having human-aware navigation for robots in 2B-SD, robots could not fully eliminate social costs, although M_3 - M_5 has been improved for 2B-SD compared to 2B-AG on average. We observed two problematic situations for the robots in term of human-aware navigation. Firstly, when Robot 1 is moving to Room 1, Human 1 is about to exit the room. The doorway is a very narrow passage that does not allow for both of them to progress. This is a hard




Fig. 5: Task assignment per robot over time for a sample run of the second set of experiments for scenarios 2A, 2B-AG, 2B-SD, and 2B-SR-KF respectively. End of mission (M_2) is marked by the vertical line.

situation for human-aware navigation and we observed that Robot 1 partially intruded the social space of the human with variable severity in different runs. Secondly, there were a few cases of collision when a human was moving too fast for the robots to be able to adjust their paths and our individual human-aware navigation was not able to ensure respecting the social constraints in more difficult and dynamic situations.

Contrarily, in Scenario 2B-SR-KF because of avoiding risky areas and despite having no individual human-aware navigation, robots were able to ensure social-awareness. Risk-based planning does not promise less traveled distance or less time (see Fig. 4a and 4b). Depending on the human behavior it might find a plan that is longer or shorter in terms of distance and time, but it ensures that decisions will be made considering all aspects (including social costs) together.

VII. CONCLUSION AND FUTURE WORK

In this paper, we have proposed a risk-based bid estimation method for MRTA with stochastic costs in dynamic social environments. We investigated the effect of prediction error on the performance of different risk formulations. Results confirm that risk-based plans that account for social costs lead to better team plans in terms of social metrics, prevent difficult social situations, and reduce the need for the lower level human-aware navigation to be activated. Although risk-based planning alone was able to achieve socially acceptable results in our test cases, the combination of risk-based human-aware coordination and planning, and human-aware individual navigation ensures that social constraints will be respected even if higher level plans incur some social costs due to yet unpredictable changes in the environment or other sources of uncertainty.

For future steps of our work, we will focus on adaptive risk-based replanning that can improve the team performance by correcting inaccurate or invalid estimates that can occur as a result of sudden changes in the environment. Additionally, we will look into MRTA with limited local perception and information sharing among robots.

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] A. D. Dragan, K. C. Lee, and S. S. Srinivasa, "Legibility and predictability of robot motion," in *8th ACM/IEEE International Conference on Human-Robot Interaction*, pp. 301–308, 2013.
- [3] B. P. Gerkey and M. J. Mataric, "A formal analysis and taxonomy of task allocation in multi-robot systems," *The International Journal of Robotics Research*, vol. 23, no. 9, pp. 939–954, 2004.
- [4] G. A. Korsah, A. Stentz, and M. B. Dias, "A comprehensive taxonomy for multi-robot task allocation," *The International Journal of Robotics Research*, vol. 32, no. 12, pp. 1495–1512, 2013.
- [5] M. B. Dias, R. Zlot, N. Kalra, and A. Stentz, "Market-based multirobot coordination: A survey and analysis," *Proceedings of the IEEE*, vol. 94, no. 7, pp. 1257–1270, 2006.
- [6] E. Nunes, M. Manner, H. Mitiche, and M. Gini, "A taxonomy for task allocation problems with temporal and ordering constraints," *Robotics and Autonomous Systems*, vol. 90, pp. 55–70, 2017.
- [7] A. Khamis, A. Hussein, and A. Elmogy, "Multi-robot task allocation: A review of the state-of-the-art," in *Cooperative Robots and Sensor Networks*, pp. 31–51, 2015.
- [8] Y. Cao, W. Yu, W. Ren, and G. Chen, "An overview of recent progress in the study of distributed multi-agent coordination," *IEEE Transactions on Industrial Informatics*, vol. 9, no. 1, pp. 427–438, 2013.
- [9] L. Xu and A. Stentz, "Market-based coordination of coupled robot systems," in *2011 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 2784–2789, IEEE, 2011.
- [10] M. G. Lagoudakis, E. Markakis, D. Kempe, P. Keskinocak, A. J. Kleywegt, S. Koenig, C. A. Tovey, A. Meyerson, and S. Jain, "Auction-based multi-robot routing," in *Robotics: Science and Systems*, vol. 5, pp. 343–350, 2005.
- [11] M. Gendreau, G. Laporte, and R. Séguin, "Stochastic vehicle routing," *European Journal of Operational Research*, vol. 88, no. 1, pp. 3–12, 1996.
- [12] S. S. Ponda, L. B. Johnson, and J. P. How, "Distributed chance-constrained task allocation for autonomous multi-agent teams," in *American Control Conference*, pp. 4528–4533, 2012.
- [13] J. F. Quindlen and J. P. How, "Machine learning for efficient sampling-based algorithms in robust multi-agent planning under uncertainty," in *AAAI Guidance, Navigation, and Control Conference*, pp. 1921–1940, 2017.
- [14] J. Capitan, M. T. Spaan, L. Merino, and A. Ollero, "Decentralized multi-robot cooperation with auctioned pomdps," *The International Journal of Robotics Research*, vol. 32, no. 6, pp. 650–671, 2013.
- [15] K. E. Booth, S. C. Mohamed, S. Rajaratnam, G. Nejat, and J. C. Beck, "Robots in retirement homes: Person search and task planning for a group of residents by a team of assistive robots," *IEEE Intelligent Systems*, vol. 32, no. 6, pp. 14–21, 2017.
- [16] K. E. Booth, G. Nejat, and J. C. Beck, "A constraint programming approach to multi-robot task allocation and scheduling in retirement homes," in *International Conference on Principles and Practice of Constraint Programming*, pp. 539–555, 2016.
- [17] P. Khandelwal and P. Stone, "Multi-robot human guidance: Human experiments and multiple concurrent requests," in *Proceedings of the 16th Conference on Autonomous Agents and Multi-Agent Systems*, pp. 1369–1377, 2017.
- [18] Z. Talebpour, S. Savarè, and A. Martinoli, "Market-based coordination in dynamic environments based on the Hoplites framework," in *The IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 1105–1112, 2017.
- [19] N. Kalra, D. Ferguson, and A. Stentz, "Hoplites: A market-based framework for planned tight coordination in multirobot teams," in *IEEE International Conference on Robotics and Automation*, pp. 1170–1177, 2005.
- [20] Z. Talebpour and A. Martinoli, "Multi-robot coordination in dynamic environments shared with humans," in *IEEE International Conference on Robotics and Automation*, 2018.
- [21] J. Rios-Martinez, A. Spalanzani, and C. Laugier, "Understanding human interaction for probabilistic autonomous navigation using risk-RRT approach," in *The IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 2014–2019, 2011.
- [22] E. T. Hall, R. L. Birdwhistell, B. Bock, P. Bohannon, A. R. Diebold Jr, M. Durbin, M. S. Edmonson, J. Fischer, D. Hymes, S. T. Kimball, et al., "Proxemics [and comments and replies]," *Current anthropology*, vol. 9, no. 2/3, pp. 83–108, 1968.
- [23] J. V. Gomez, N. Mavridis, and S. Garrido, "Social path planning: Generic human-robot interaction framework for robotic navigation tasks," in *2nd Intl. Workshop on Cognitive Robotics Systems: Replicating Human Actions and Activities*, 2013.
- [24] R. Ventura and A. Ahmad, "Towards optimal robot navigation in domestic spaces," in *RoboCup 2014: Robot World Cup XVIII*, pp. 318–331, Springer, 2015.
- [25] J. V. Gómez, N. Mavridis, and S. Garrido, "Fast marching solution for the social path planning problem," in *IEEE International Conference on Robotics and Automation*, pp. 2243–2248, 2014.
- [26] Z. Talebpour, D. Viswanathan, R. Ventura, G. Englebienne, and A. Martinoli, "Incorporating perception uncertainty in human-aware navigation: A comparative study," in *25th International Symposium on Robot and Human Interactive Communication*, pp. 570–577, 2016.
- [27] J. Sequeira, P. Lima, A. Saffioti, V. Gonzalez-Pacheco, and M. Salichs, "MONarCH: Multi-robot cognitive systems operating in hospitals," in *ICRA workshop on many robot systems*, 2013.
- [28] J. Messias, R. Ventura, P. Lima, J. Sequeira, P. Alvaro, C. Marques, and P. Carriço, "A robotic platform for edutainment activities in a pediatric hospital," in *IEEE International Conference on Autonomous Robot Systems and Competitions*, pp. 193–198, 2014.
- [29] O. Michel, "Webots: Symbiosis between virtual and real mobile robots," in *Virtual Worlds*, pp. 254–263, 1998.