

A Semantically-Informed Multirobot System for Exploration of Relevant Areas in Search and Rescue Settings

Alberto Quattrini Li · Riccardo Cipolleschi · Michele Giusto · Francesco Amigoni

Received: date / Accepted: date

Abstract Coordinated multirobot exploration involves autonomous discovering and mapping of the features of initially unknown environments by using multiple robots. Autonomously exploring mobile robots are usually driven, both in selecting locations to visit and in assigning them to robots, by knowledge of the already explored portions of the environment, often represented in a metric map. In the literature, some works addressed the use of *semantic knowledge* in exploration, which, embedded in a semantic map, associates spatial concepts (like ‘rooms’ and ‘corridors’) with metric entities, showing its effectiveness in improving the *total* area explored by robots. In this paper, we build on these results and propose a system that exploits semantic information to push robots to explore *relevant* areas of initially unknown environments, according to *a priori* information provided by human users. Discovery of relevant areas is significant in some search and rescue settings, in which human rescuers can instruct robots to search for victims in specific areas, for example in cubicles if a disaster happened in an office building during working

hours. We propose to speed up the exploration of specific areas by using semantic information both to select locations to visit and to determine the number of robots to allocate to those locations. In this way, for example, more robots could be assigned to a candidate location in a corridor, so the attached rooms can be explored faster. We tested our semantic-based multirobot exploration system within a reliable robot simulator and we evaluated its performance in realistic search and rescue indoor settings with respect to state-of-the-art approaches.

Keywords Coordinated multirobot exploration · Semantic map · Search and rescue

1 Introduction

Coordinated multirobot exploration (Burgard et al, 2005) autonomously discovers features of initially unknown environments by using mobile robots equipped with sensors. Exploration is fundamental in tasks like map building (Thrun, 2002) and search and rescue (Tadokoro, 2010). Decisions about where to go next and about which robot goes where are crucial in coordinated multirobot exploration and are usually made according to information extracted from the known portion of the environment, represented in a *metric map* that robots incrementally build. A metric map represents the spatial features of the environment, like the position of obstacles. In the last years, several methods have been proposed to build *semantic maps* of environments (like (Wolf and Sukhatme, 2008) and (Mozos et al, 2005)), which label some spatial elements with high-level human concepts. For example, areas of a metric map can be labeled as ‘corridor’ or ‘room’, thus providing knowledge about the structure of the environment. Despite

Alberto Quattrini Li (✉) · Riccardo Cipolleschi · Michele Giusto · Francesco Amigoni
Artificial Intelligence and Robotics Laboratory
Dipartimento di Elettronica, Informazione e Bioingegneria,
Politecnico di Milano
Piazza Leonardo da Vinci 32, 20133
Tel.: +39 02 2399 9687
Fax: +39 02 2399 3411
E-mail: alberto.quattrini@polimi.it

Riccardo Cipolleschi
E-mail: riccardo.cipolleschi@mail.polimi.it

Michele Giusto
E-mail: michele.giusto@mail.polimi.it

Francesco Amigoni
E-mail: francesco.amigoni@polimi.it

the great effort in *constructing* semantic maps, the study of their *use* for exploration is still rather limited.

In this paper, we contribute to this study by presenting a coordinated multirobot exploration system that operates in search and rescue settings and that exploits semantic labels to explore *relevant* areas of environments. A relevant area is defined as a portion of the environment that is considered of interest according to what human users communicate to the system. Our system is composed of multiple robots equipped with laser range scanners that operate according to the following steps: they (a) perceive the surrounding environment, (b) integrate the perceived data within a metric and a semantic map representing the environment known so far, (c) decide where to go next and who goes where, and (d) go to the assigned target locations and start again from (a). In this work, we focus on the decision making (step (c)).

Some works (e.g., (Calisi et al, 2009; Stachniss et al, 2008)) have already addressed the problem of exploiting semantic knowledge to improve exploration, finding out that the use of semantic information can reduce the time required to cover a given amount of area and can increase the *total* amount of area mapped by robots in a given time interval. In this paper, we extend these results by showing that semantic knowledge can also be used to significantly improve the exploration of relevant areas of indoor environments. We assume that *a priori* and reliable information about the areas of the environment that are considered relevant is available, for example, provided by humans. This assumption is of interest in realistic scenarios. In a search and rescue setting, the *a priori* information could be the possible location of victims or the preferred areas to search first, given by human rescuers. For example, if a disaster happens during office hours, victims are most likely located in the offices, and, thus, robots should focus on searching small-size rooms. If it happens during lunch time, robots should head to large-size rooms, like a canteen. In the following of this paper, we consider two kinds of *a priori* information about victim location: victims in small rooms and victims in big rooms. However, our approach can be applied to more general settings with arbitrary *a priori* information coming from human users. We propose to exploit semantic information to select interesting locations to visit and, differently from the literature, to assign more robots to a single location. For example, in an indoor environment, if a location lies in an area labeled as ‘corridor’, then that area could be privileged by sending there several robots, so that rooms, typically attached to corridors, can be explored faster.

Our system originally addresses the following problem: To what extent is it possible and convenient to exploit semantic information to efficiently explore areas (of an initially unknown environment) that are considered relevant? To the best of our knowledge, none of the works present in the literature has addressed such research question. The main original contribution of this paper is thus a system that exploits semantic information to improve exploration of relevant areas. Specifically, we propose a method for evaluating candidate locations, which is a variant of that presented in (Basilico and Amigoni, 2011), and a method for allocating robots to candidate locations. The contributions of this paper significantly extend the preliminary results in (Cipolleschi et al, 2013) and include a more complete experimental analysis of the behavior and benefits of the proposed multirobot system, also involving additional experiments using state-of-the-art approaches for comparison and the deployment of more robots.

This paper is structured as follows. The next section surveys the methods that have been proposed to perform the decision making step (c) in the context of coordinated multirobot exploration. Section 3 presents the proposed multirobot exploration system. Section 4 shows extensive experimental simulated activities with the aim to display the effectiveness of the proposed approach. Finally, Section 5 concludes the paper.

2 Related work

Robotic exploration is the task in which mobile robots, equipped with on-board sensors, are employed in the iterative online process presented in the previous section with the goal to discover (initially unknown) features in environments. The mainstream approach to robotic exploration (Yamauchi, 1998) identifies some candidate locations on the *frontiers* between known (already explored) and unknown portions of the environment, evaluates them, and assigns them to robots, iteratively. In this paper, we are interested in problems of searching targets in an initially unknown environment, by using multiple mobile robots, with a semantic map and some *a priori* reliable information about the location of the targets (e.g., victims most likely located in offices of a building), which is provided by humans. In addressing these issues, and in order to clearly present our contribution, we tear apart the aspects of evaluating the candidate locations (*exploration strategy*) and of allocating robots to candidate locations (*coordination method*). In the following, we present a representative sample of the several exploration strategies and coordination methods (focusing on those using semantic information

for speeding up autonomous exploration of initially unknown environments) presented in literature.

2.1 Exploration strategies

In the literature, exploration strategies usually employ a utility function that combines different criteria, which characterize each candidate location, to assess the goodness of different candidate locations. Most of the criteria considered by exploration strategies are only relative to metric information, namely information that can be derived from metric maps that robots build. For example, the robotic exploration system of González-Baños and Latombe (2002) combines, in an exponential function, the distance between a robot r and a candidate location p and the expected amount of information that r can acquire at p (measured as the maximum amount of unknown area visible from p). A system using the same two criteria, but combining them in a linear function, is that of Burgard et al (2005). The system proposed by Basilico and Amigoni (2011) adds a criterion that measures the probability that r , once in p , can communicate with a fixed base station (like Visser and Slamet (2008)), and combines all the criteria using a theoretically-grounded approach.

Only few exploration systems use semantic information to evaluate candidate locations and assign them to the robots. An early attempt in this direction is that of Kuipers and Byun (1981), in which candidate locations with a large distinctiveness (e.g., located at the intersections of corridors) are privileged. Specifically, the authors, without explicitly considering semantic information, use a geometric measure, which derives from finding points that are equally-distant to close obstacles, and apply a hill-climbing control strategy to find the robot’s exploration path, so that distinctiveness is maximized. They show the goodness of the exploration path found through a qualitative analysis of the solution obtained during simulations.

The work in (Stachniss et al, 2008) exploits the knowledge on the structure of an indoor environment (represented as a hidden Markov model) to drive robots to select, first, candidate locations that are in corridors. For each robot r and each candidate location p , the difference between the initial utility of p (which is equal for all frontiers and is initialized at 1, not considering any features of p) and the distance between the current position of r and p is calculated. The initial utility of candidate locations that are in corridors is multiplied by γ (set to 5 in the experimental activity). Then the method greedily allocates the candidate locations to the robots by selecting the pair r and p that maximizes the above difference. Experimental results (performed

in simulation) show that the approach is effective in decreasing the time required to explore some environments with respect to an approach that does not update the initial utility of corridors.

Another work that uses semantic information to improve exploration is presented in (Calisi et al, 2009). In this case, contextual information related to the mission (e.g., the relative importance of a goal with respect to another goal), to the environment (e.g., the presence of rooms and corridors and the difficulty for traversing a given area and for detecting victims in that area), and to the agents (e.g., the presence of loop closures for improving localization of robots) is represented by a PROLOG rule-based system and exploited to enhance the performance of a robotic system operating in a search and rescue scenario. The experiments (performed in simulation) use a single robot and show that the proposed approach can significantly increase the area mapped by the robot within 15 minutes.

Another system that exploits the structure of the environment for determining the best candidate locations and assigning them to the robots is presented in (Wurm et al, 2008). The known portion of the map of the environment is segmented and a single robot is assigned to (one of the frontiers of) each segment. The utility function used to assign a robot r to a frontier p considers the distance from the current position of r to p . Experimental results in simulation show that the approach can significantly reduce the overall exploration time for realistic environments with respect to a closest-frontier approach that assigns to each robot the closest candidate location. Also, the authors validated the method with two real robots in an environment, qualitatively analyzing the paths they follow.

All these works that embed semantic information in the exploration strategy show that the *total* area explored in a given time interval can be improved by using semantic information, but do not exploit semantic information to push robots to explore areas that are *relevant*, according to *a priori* information available from, for example, human rescuers. Moreover, all these approaches use a utility function that considers basically just the cost to reach a candidate location. Amigoni (2008) experimentally showed that, in some common settings, exploration strategies that balance utility and cost tend to have better performance than those that use only cost.

Since in this paper we are assuming that some *a priori* information is available about relevant areas, the above works are not directly comparable with our approach, in terms of the relevant areas explored. However, in our experimental simulated activities, we will compare the proposed exploration strategy with that

proposed by Basilico and Amigoni (2011), which has been experimentally proven to perform well in search and rescue scenarios, but do not consider any semantic information. In this paper, we aim at showing that, when *a priori* knowledge on victims' locations is available (i.e., preferred areas to visit are specified), the use of semantic information could improve also the performance of exploration of relevant areas of the environment, besides the total explored area, as shown by the works presented above (like (Stachniss et al, 2008) or (Calisi et al, 2009)).

2.2 Coordination methods

Coordination between multiple exploring robots, namely assigning robots to candidate locations, is achieved in different ways in the literature.

A series of works ((Burgard et al, 2000, 2005; Stachniss et al, 2008) and, partially, (Fox et al, 2006)) propose an interesting approach in which the coordination method is embedded within the exploration strategy. In particular, the utility value of a candidate location is reduced according to the number of robots that can view it, in order to discourage the assignment of more robots to the same candidate location. Experimental results show that this coordinated behavior has better performance than uncoordinated behavior (in which different robots can select the same location to reach) and slightly worse performance than a method that finds the optimal allocation over all possible permutations of candidate locations to robots, where the optimality criterion depends on the difference between utility and cost of visiting the candidate locations.

Several other works (e.g., (Simmons et al, 2000; Zlot et al, 2002)) are based on market mechanisms. Specifically, coordination of mobile robots is performed by a central executive that, beyond collecting local maps and combining them into a single global map, manages an auction-like mechanism by asking bids to the robots and assigning tasks (i.e., locations to reach) according to the received bids. Bids contain information about expected utility for pairs robot-location; utility are calculated by the exploration strategy adopted. Experimental results show that the auction-based coordination methods (as expected) outperform the uncoordinated methods. An extension of such works is the approach of Hawley and Butler (2013), who propose an auction-based coordination method not only for task assignment, but also for coalition formation, when there are more robots than candidate locations.

All of the presented approaches for coordination attempt to spread the robots around the environment. The (often) implicit assumption is that the exploration

problem is considered to involve, according to the classification in (Gerkey and Mataric, 2004), single-task robots (ST) and single-robot tasks (SR), where the task is to reach a candidate location. ST means that each robot executes one task at a time and SR means that each task requires one robot. Thus, all the above works act basically as ST-SR.

In this paper, we attempt to overcome the ST-SR assumption by allocating more robots to the same candidate locations according to a multi-robot tasks (MR) paradigm. Specifically, we aim at showing that semantic information enables the possibility to determine the ideal number of robots to send to a specific area so that exploration can proceed faster and more effectively.

3 A Semantic-based multirobot exploration system

In this section, we present our proposed exploration system that exploits semantic information. Specifically, after a brief system overview, we go into the details of the exploration strategy and the coordination method we designed for the semantic-based multirobot exploration system.

3.1 System overview

The robotic platform used is a Pioneer 3-DX equipped with a sonar ring and two laser range scanners, mounted at the same height and back-to-back for covering a 360° area around the robot with radius $R = 20$ m and angular resolution at 1°.

Each robot builds a two-dimensional occupancy grid map of the explored environment. Each cell is either known, if the robot perceived the corresponding area, or unknown. Known cells can be free or occupied (by obstacles). The map of the environment is maintained by a base station, whose position is fixed in the environment, and to which robots send their maps every 2.5 s. We assume that communication is error-free and unlimited in range and bandwidth (effects of more realistic communication models on exploration are discussed in (Tuna et al, 2012)). Our exploration system is largely independent of the mapping system employed to incrementally build the grid map. In our experiments, we use a simple scan matching method, inspired to that in (Lu and Milios, 1997), in which a new acquired scan is aligned with the current map (using odometry as initial guess) and the occupancy grid is updated correspondingly. Since we are not interested in analyzing the quality of the resulting map, we assume that the mapping module is error-free. Given the grid map, clusters

of (adjacent) free cells that are on the frontier between known and unknown parts of the map are extracted. For each cluster, the free cell belonging to that cluster and closest to its centroid is considered as a *candidate location* to reach. Paths are planned using A* on the grid map. Sonars are used for obstacle detection during navigation.

We assume that the system has a semantic map that labels each free cell of the grid map with its room type (i.e., labels ‘corridor’, ‘small room’, ‘medium room’, ‘big room’) and with the number of doorways present in the room in which the cell is located. This semantic map can be built exploiting any available method (e.g., (Mozos et al, 2005)). However, in this paper we assume the semantic map as available, because we are only interested in its use. In practice, we manually annotate with semantic labels the portions of the simulated environments used for the experimental activities. Note that the proposed approach can be, in principle, applied to any number of semantic labels, different from the four we consider.

3.2 MCDM-based exploration strategy

Our Multi-Criteria Decision Making (MCDM) exploration strategy uses several criteria to evaluate the goodness of a candidate location. More formally, the exploration strategy is used to estimate the utility $u(p, r)$ of every candidate location p for all robots r . It combines the following criteria:

- $A(p)$ is the expected amount of free area beyond the frontier of p computed as the length (in cells) of the frontier. The larger its value, the more information is expected to be acquired from p .
- $d(p, r)$ is the Euclidean distance between p and current position of r . Using Euclidean distance instead of actual distance calculated by path planner drastically reduces the computational effort in calculating this criterion without affecting too much the estimated utility $u(p, r)$, as some preliminary experiments we performed have shown.
- $b(p, r)$ is an estimate of the energy spent by r for reaching p , calculated considering a very simple model, in which the power consumption is related to the time required for reaching p , computed according to the path that r should follow and according to linear and angular velocities of the robots. The larger its value, the smaller the amount of residual energy in the battery (0 = full, 1 = empty).

All these criteria can be calculated from the robots’ status and from the metric grid map.

In addition to the above criteria, other criteria employing information from semantic map are considered:

- $S(p)$ is the relevance of p (from 0, not relevant, to 1, relevant), calculated according to the semantic label of p and the *a priori* knowledge on victims’ locations. For example, if it is known that victims are most likely in big rooms, and p is labeled as ‘big room’, $S(p) = 1$, while if p is labeled as ‘small room’, and under the same hypothesis about the location of victims, $S(p) = 0$. If p is labeled as ‘corridor’, regardless the hypothesis on victims’ locations, $S(p) = 0.15$, as corridors are usually important to reach relevant rooms. The values for $S(p)$ have been manually set to obtain good performance after experiments with different combinations of values. In our preliminary tests, different value combinations (e.g., range [0.10, 0.50] for $S(p)$ with p in corridors), that maintain relevance of corridors and of rooms according to the hypothesis on victims’ location, have been experimentally demonstrated to have similar performance.
- $ND(p)$ is the number of doors in the room where p is located. This criterion evaluates the connectivity of a room with other rooms. The idea is that a highly-connected room should be visited to ease finding relevant rooms.

We assume that semantic labeling used to calculate the criteria $S(p)$ and $ND(p)$ is perfect. This assumption will be relaxed later to experimentally verify the robustness of the approach. Note that, in order to apply our approach to other semantic labels and other kinds of *a priori* information, criterion $S()$ should be changed.

All the criteria $N = \{A, d, b, S, ND\}$ are combined using the MCDM approach introduced in (Basilico and Amigoni, 2011), to which we refer for a complete description; here we just summarize the main features. We selected the MCDM approach because it is theoretically grounded and allows to easily integrate several criteria in a utility function. Consider a set of candidate locations \mathcal{P} (i.e., the cells closest to the centroids of their frontiers at some time during exploration), a set of robots \mathcal{R} , and a set of criteria N . Call $u_j(p, r)$ the utility value for candidate location $p \in \mathcal{P}$ and robot $r \in \mathcal{R}$ according to criterion $j \in N$. The larger $u_j(p, r)$, the better the pair p and r according to j . To apply MCDM, utilities need to be normalized to a common scale $I = [0, 1]$. We use a linear relative normalization for each u_j . With a slight abuse of notation, we call $u_{(j)}$, with $(j) \in N$, the j -th criterion according to an increasing ordering with respect to utilities: for candidate location p and robot r , $u_{(1)}(p, r) \leq \dots \leq u_{(n)}(p, r) \leq 1$, where $n = |N|$ (we assume $u_{(0)}(p, r) = 0$). The MCDM

S-MCDM	criteria	$\mu()$	criteria	$\mu()$	criteria	$\mu()$
	A	0.09	d, b	0.09	A, b	0.15
	d	0.09	d, S	0.8	d, b, S	0.8
	b	0.02	b, S	0.6	A, S	0.65
	S	0.5	A, d, b	0.3	A, d, b, S	0.8
	ND	0.3	A, d, S	0.8		
	A, d	0.3	A, b, S	0.65		

D-MCDM	criteria	$\mu()$
	A	0.4
	d	0.4
	b	0.2
	A, d	0.95
A, b	0.7	
d, b	0.4	

Table 1 Weights of MCDM-based exploration strategies.

strategy integrates the criteria in N with the following function:

$$u(p, r) = \sum_{j=1}^n (u_{(j)}(p, r) - u_{(j-1)}(p, r)) \mu(\mathcal{A}_{(j)}), \quad (1)$$

where $\mu : 2^N \rightarrow [0, 1]$ (2^N is the power set of set N) are weights, and the set $\mathcal{A}_{(j)}$ is defined as $\mathcal{A}_{(j)} = \{i \in N | u_{(j)}(p, r) \leq u_i(p, r) \leq u_{(n)}(p, r)\}$. Specifically, $\mu(\{\emptyset\}) = 0$, $\mu(N) = 1$, and, if $N' \subset N'' \subset N$, then $\mu(N') \leq \mu(N'')$. That is, μ is a normalized *fuzzy measure* on the set of criteria N that will be used to associate a weight to each group of criteria. The weights specified by the definition of μ describe the relationships between criteria. Criteria belonging to a group $G \subseteq N$ are said to be redundant if $\mu(G) < \sum_{i \in G} \mu(i)$, synergic if $\mu(G) > \sum_{i \in G} \mu(i)$, and independent otherwise. Namely, Equation (1) provides a sort of “distorted” weighted average that accounts for synergies and redundancies between criteria.

In MCDM, beyond selecting the set of criteria N , we need to define weights μ for each subset of criteria. For our semantically-informed exploration strategy (*S-MCDM*), we use the criteria $N = \{A, d, b, S, ND\}$ defined above and the weights reported in Table 1 (left). The weights of the subsets of criteria not reported in the table are calculated by summing the weights of the individual criteria. Note that in selecting these weights, we have chosen values reasonably (e.g., criteria $d()$ and $A()$ have the same importance, so their weights are equal). Moreover, criteria $d()$ and $b()$ are redundant (both prefer candidate locations close to the robot and a candidate location satisfies both criteria well or both not well) and so $\mu(\{d, b\}) < \mu(\{d\}) + \mu(\{b\})$. Criteria $A()$ and $d()$ are instead synergic (one prefers candidate locations on long frontiers while the other one prefers candidate locations close to the robot and a candidate location can satisfy one criterion well and the other one not well) and so $\mu(\{A, d\}) > \mu(\{A\}) + \mu(\{d\})$. Values of weights have been set to obtain good performance, according to criteria importance and relations (Basilico and Amigoni, 2011). Slightly varying the selected weights values ($\pm 10\%$), we experimentally obtained similar performance. Principled methods for selecting weights are discussed in (Basilico and Amigoni, 2011).

For comparing the performance of S-MCDM, we chose a state-of-the-art exploration strategy. Specifically, we defined another MCDM-based exploration strategy (called in the following Default MCDM, or *D-MCDM*), whose criteria set is $N = \{A, d, b\}$, similarly to (Basilico and Amigoni, 2011), and with weights reported in Table 1 (right). As discussed in Section 2 and to the best of our knowledge, no exploration strategy that focuses on relevant areas is available. Furthermore, the work of Calisi et al (2009) is not easily configurable in our setting, as PROLOG rules should be set. Nevertheless, the D-MCDM exploration strategy has been shown by Basilico and Amigoni (2011) to be very effective in exploring environments (in particular, it outperformed the exploration strategies proposed by Visser and Slamet (2008) and Amigoni and Caglioti (2010)).

3.3 ST-MR coordination method

Coordination methods are used to assign candidate locations to robots. The mechanism we use is market-based (Zlot et al, 2002). The base station regularly sets up auctions in which candidate locations (generated on current frontiers as discussed before) are auctioned to the robots, which bid on them. This process allocates candidate locations p to robots r attempting to maximize the sum of utilities $u(p, r)$. In our system, the coordination method can allocate multiple robots (MR) to the same candidate locations. For example, allocating two robots to the same candidate location in a big room could speed up the exploration of the room, overcoming potential negative effects due to the initially overlapping views of the two robots.

We employ a fuzzy-based function $i(p)$ that computes the ideal number of robots (1, 2, or 3, in our experiments) that should be assigned to a candidate location p , according to the semantic label given to p and to some other features. In particular, if p is located in a room (‘small room’, ‘medium room’, or ‘big room’), the features considered are the room area, the free area percentage of the total area in the room (visibility), the number of doors, and the already perceived area of the room. Note that an estimate of the already perceived area of a room can be computed by having a knowledge base that associates the semantic labels of rooms to the corresponding average area (see, e.g., the work of Luperto et al (2013)). Fig. 1 illustrates the membership functions for the input features and for the output for p in a room that we have used for experiments we show in the next section. When slightly varying the selected fuzzy values ($\pm 10\%$), we experimentally obtained similar performance. Given p , if the room in which p is located is large, the number of its doors is large, its

visibility is large, and the amount of already perceived area is small, then more robots are allocated to p . Another example of the rules for determining the ideal number of robots $i(p)$ to be allocated to p (in a room) is reported in Algorithm 1.

```

1 if RoomSize is SMALL and #Doors is HIGH and Visibility
  is LOW and AlreadyPercArea is MEDIUM then #Robots
  is MEDIUM;
2 if RoomSize is BIG and #Doors is LOW and Visibility is
  LOW and AlreadyPercArea is HIGH then #Robots is
  LOW;
3 if RoomSize is BIG and #Doors is MEDIUM and
  Visibility is MEDIUM and AlreadyPercArea is LOW then
  #Robots is HIGH;

```

Algorithm 1: Sample of rules for calculating the ideal number of robots that can be allocated to p (in a room).

Similarly, if p is located in a corridor (label ‘corridor’), the features considered are the size of the corridor, the number of doors, the number of intersecting corridors, and the already perceived area of the corridor. The membership functions and the rules are similar to those for the room case, as shown in Fig. 2 and in Algorithm 2. Note that, in order to use our approach with different semantic labels, membership functions and rules for calculating $i(p)$ should be changed.

```

1 if CorridorSize is SMALL and #Doors is HIGH and
  #IntersectingCorridors is MEDIUM and AlreadyPercArea
  is MEDIUM then #Robots is LOW;
2 if CorridorSize is SMALL and #Doors is MEDIUM and
  #IntersectingCorridors is LOW and AlreadyPercArea is
  LOW then #Robots is MEDIUM;
3 if CorridorSize is MEDIUM and #Doors is MEDIUM and
  #IntersectingCorridors is MEDIUM and AlreadyPercArea
  is LOW then #Robots is HIGH;
4 if CorridorSize is BIG and #Doors is HIGH and
  #IntersectingCorridors is MEDIUM and AlreadyPercArea
  is LOW then #Robots is VERY_HIGH;

```

Algorithm 2: Sample of rules for calculating the ideal number of robots that can be allocated to p (in a corridor).

Each robot r evaluates all candidate locations p , as auctioned by the base station every 5 s or when requested by a robot that has reached its assigned location, according to the exploration strategy, and submits bids $u(p, r)$ accordingly. We propose two coordination methods executed by the base station to allocate candidate locations to robots. The first coordination method (*MRv1*) works as reported in Algorithm 3. Basically, MRv1 greedily allocates the best pair (p^*, r^*) , avoiding to allocate p^* to more than $i(p^*)$ robots.

The second coordination method, called *MRv2*, is similar to MRv1, but, after each allocation of a robot

```

1 collect bids  $u(p, r)$ , which are calculated using (1);
2 while  $\exists$  robot  $r$  not allocated and candidate location  $p$  do
3   find the pair  $(p^*, r^*)$ :  $(p^*, r^*) = \arg \max_{p, r} u(p, r)$ ;
4   allocate  $p^*$  to  $r^*$ ;
5   if  $i(p^*)$  is equal to the number of robots already assigned
     to  $p^*$  then
6     | eliminate  $p^*$ ;
     end
7   eliminate robot  $r^*$ ;
end

```

Algorithm 3: MRv1.




Fig. 3 Discount factor vs. the number of robots already allocated to p , when $i(p) = 3$.

to a p^* (step 4), it discounts the utility of p^* for other robots, according to the number of robots already allocated to p^* (similarly to (Stachniss et al, 2008)). Fig. 3 shows the discount factor we employ that decreases linearly until the number of allocated robots is less than or equal to $i(p^*)$, and then decays exponentially. The rationale is that assigning to p^* less robots than $i(p^*)$ could be a necessity (e.g., there are not enough robots) and that assigning to p^* more robots than $i(p^*)$ is not useful to speed up exploration.

The two proposed ST-MR coordination methods are experimentally compared to a standard coordination method (ST-SR) (Zlot et al, 2002), which allocates just one robot to a candidate location in a greedy fashion. Namely, it runs MRv1 with $i(p) = 1$ for every p .

4 Experimental activity

This section, first, shows the experimental setup in which we tested our proposed semantic-based exploration system. Then, we show some preliminary experiments to support the choice of the state-of-the-art exploration strategy against which our system is compared, and we present extensive experimental results that validate the system. Further, we present additional experiments for showing the robustness of our proposed system, for example by relaxing the assumption on the perfect semantic knowledge and by adopting a different termination criterion. Finally, we discuss the obtained results.

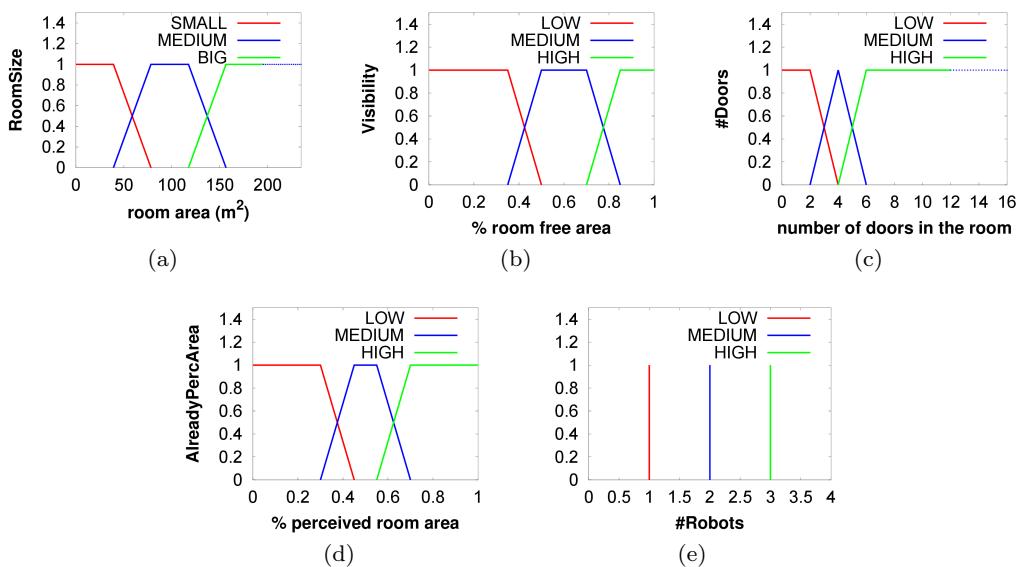


Fig. 1 Membership functions for the input features (a-d) and for the output (e), when p is in a room.

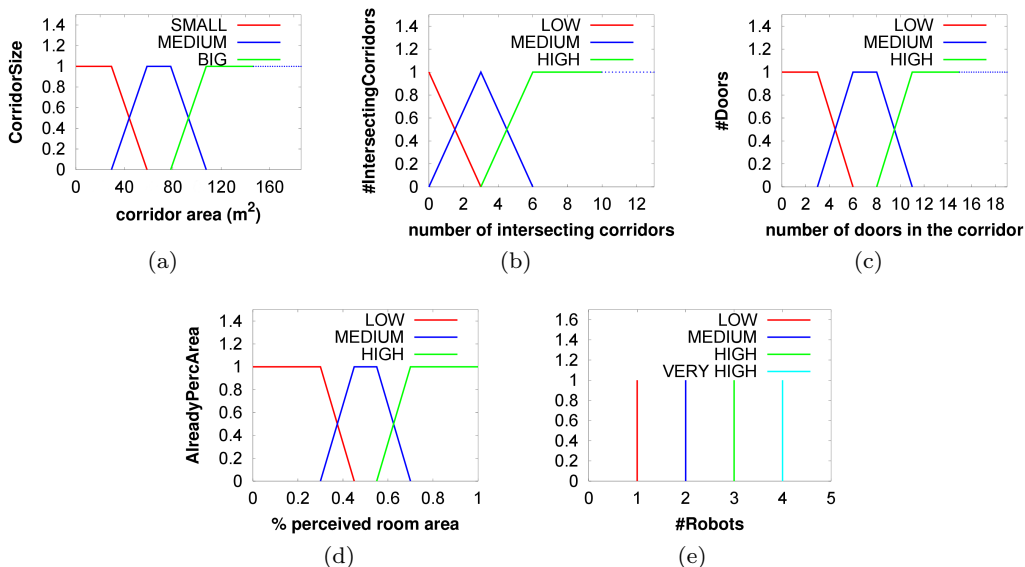


Fig. 2 Membership functions for the input features (a-d) and for the output (e), when p is in a corridor.

4.1 Experimental setup

In order to perform replicable tests under controlled conditions, we use a robot simulator. We selected US-ARSim (Carpin et al, 2007), because it is a realistic and reliable 3D robot simulator. The multirobot system controller software we developed and the experimental data are publicly available at <http://sourceforge.net/projects/polimirobotcup>.

We report simulated experiments conducted in two indoor environments, called *office* and *mall* (Fig. 4), where robots start from fixed starting locations without any initial knowledge about the structure of the

environment. The cells of the test environments are manually labeled as ‘corridor’, ‘small room’, ‘medium room’, or ‘big room’ according to the size of the rooms they belong to. Label distributions are reported in Tables 2 and 3. The office environment is part of the *vasche.library.floor1* taken by Radish repository (Howard and Roy, 2003), and is characterized mainly by the presence of small and medium rooms (as we can see from Table 2, the number of small and medium rooms is almost the 86% of the total number of rooms in the environment). The mall environment is a floor of a (real) mall, and is characterized by the presence of very big




Fig. 4 Test environments. Green stars represent initial positions for the robots in the configurations with 4 robots, red crosses refer to the addition of two robots (6 robots), and blue points to the addition of further two robots (8 robots).

Type	Number of cells	% of the environment area	Number of rooms
Corridors	2493	30%	-
Small rooms	756	9%	21
Medium rooms	2835	34%	42
Big rooms	2304	27%	10

Table 2 Number of cells, percentage of the area of the environment, and number of rooms of each semantic label (room type) for the office environment.

Type	Number of cells	% of the environment area	Number of rooms
Corridors	1449	18%	-
Small rooms	1332	16%	37
Medium rooms	2088	25%	31
Big rooms	3393	41%	9

Table 3 Number of cells, percentage of the area of the environment, and number of rooms of each semantic label (room type) for the mall environment.

rooms. Table 3 shows that the number of big rooms is almost the 12% of the total number of rooms in the environment, but they occupy 41% of the total area of the environment. Some obstacles (shown as short line segments in Fig. 4) have been added to the rooms to make the exploration task more difficult. We consider structured indoor environments because many semantic maps have been built for indoor environments and search and rescue scenarios are often indoor (like those of the Virtual Robot Competition of the RoboCup Rescue Simulation League).

We consider teams of 4, 6, and 8 robots and two *a priori* hypotheses (assumed to be correct) on victims' location, namely victims in big rooms and victims in small rooms. We define a configuration as an environment (office or mall), a number of robots (4, 6, or 8), an exploration strategy (the state-of-the-art exploration strategy D-MCDM or our proposed semantic-based exploration strategy S-MCDM, as described in Section 3.2), a coordination method (the state-of-the-art coordination method SR or our proposed coordination methods MRv1 or MRv2, shown in Section 3.3),

and an hypothesis on the victims' location (in big or small rooms). For each configuration, we execute 10 runs of 20 minutes each.

In a search and rescue setting, the goal is to explore an initially unknown environment for finding the largest number of human victims within a short time. Assuming *a priori* knowledge about the relevant area in which victims are supposed to be, and assuming that victims are uniformly distributed in such relevant areas, the problem of maximizing the number of victims found in a given time interval is equivalent to the problem of maximizing the *amount of relevant area* covered by robots' sensors in the same interval. Thus, we assess our system performance by measuring the amount of relevant area (area of small or of big rooms, according to the victims' location hypothesis) explored, every 1 minute of exploration. We typically report data at the end of runs (after 20 minutes), but, for some configurations, we report graphs of data over 20 minutes. This measure is particularly relevant in the context of search and rescue, as time is limited, and we want to explore as quickly as possible the relevant parts of an environment. We report also some results about the total explored area so that it is possible to compare our proposed method with other approaches that do not consider relevant area.

4.2 Preliminary experiments

We start with some preliminary experiments that support our choice of D-MCDM as representative state-of-the-art exploration strategy. In particular, we compare the state-of-the-art exploration strategy D-MCDM with other two exploration strategies, namely, a random one (Random, which selects the next candidate location at random) and one that only minimizes the distance (Distance), as in (Wurm et al, 2008). In all cases, the co-

Office		Exploration			
	Coord.	D-MCDM (B)	S-MCDM (B)	D-MCDM (S)	S-MCDM (S)
#Robots = 4	SR	747.3(112.5)	1473.9(202.7)	80.4(22.4)	276.2(83.9)
	MRv1	953.8(111.4)	1729.9(81.9)	103.1(22.7)	224.6(73.1)
	MRv2	921.6(130.3)	173.8(65.4)	112.9(24.4)	272.1(52.1)
#Robots = 6	SR	1024.6(220.7)	1603.2(59.1)	127.6(38.5)	235.7(43.1)
	MRv1	1123.6(143.6)	1851.3(7.8)	148.6(21.2)	400.5(100.4)
	MRv2	1164.9(94.2)	1856.7(36.3)	162.5(22.2)	429.0(83.8)
#Robots = 8	SR	1222.6(133.0)	1653.5(129.2)	170.3(43.0)	312.0(28.1)
	MRv1	1379.8(122.8)	1877.6(62.6)	186.2(31.7)	496.3(101.8)
	MRv2	1284.9(144.5)	1854.3(80.3)	185.6(42.7)	454.3(125.2)

Table 4 Results (average and standard deviation) of explored relevant area (m^2) for the office environment, after 20 minutes of exploration. B indicates victims most likely are in big rooms, S in small rooms.

ordination method is the most used in the state of the art, namely SR. We found, according to Basilico and Amigoni (2011), who found the same outcome for other environments, that D-MCDM performs better than the other two exploration strategies. For example, Fig. 5 shows that, in the case of office environment, 6 robots, SR coordination method, D-MCDM outperforms Random and performs relatively better than Distance, in terms of total explored area (measured in m^2). This provides a justification of the choice of using D-MCDM as baseline exploration strategy for comparing our proposed exploration strategy.




Fig. 5 Total explored area (m^2) over 20 minutes, in office environment, by 6 robots, with Random, Distance, and D-MCDM exploration strategies and SR coordination method.

4.3 Results for the office and the mall environments

Table 4 reports experimental results for the office environment. The values reported in each entry are the average and the standard deviation (in parentheses) over the 10 runs of the corresponding configuration.

With all the three coordination methods, our proposed semantic-based exploration strategy S-MCDM performs better than the state-of-the-art exploration strategy D-MCDM, and differences are statistically significant, according to an ANOVA analysis with a threshold for significance p -value < 0.05 (Pestman, 1998).

Office		Exploration			
	Coord.	D-MCDM (B)	S-MCDM (B)	D-MCDM (S)	S-MCDM (S)
#Robots = 4	SR	2372.7(240.1)	3424.4(316.1)	2372.7(240.1)	3956.6(556.5)
	MRv1	2753.4(199.9)	3305.3(268.9)	2753.4(199.9)	3667.1(597.4)
	MRv2	2855.0(285.2)	3616.0(219.4)	2855.0(285.2)	3683.0(441.0)
#Robots = 6	SR	3060.3(408.1)	3958.0(187.9)	3060.3(408.1)	3895.1(528.1)
	MRv1	3361.4(339.3)	4333.1(187.0)	3361.4(339.3)	5094.0(194.8)
	MRv2	3536.0(158.2)	4408.0(282.2)	3536.0(158.2)	5200.4(278.9)
#Robots = 8	SR	3612.8(324.0)	4350.6(313.5)	3612.8(324.0)	4677.6(526.5)
	MRv1	3996.5(267.9)	4856.2(524.0)	3996.5(267.9)	5424.5(264.6)
	MRv2	3853.8(358.2)	4786.8(486.7)	3853.8(358.2)	5251.2(528.2)

Table 5 Results (average and standard deviation) of total explored area (m^2) for the office environment, after 20 minutes of exploration. B indicates victims most likely are in big rooms, S in small rooms.

Mall		Exploration			
	Coord.	D-MCDM (B)	S-MCDM (B)	D-MCDM (S)	S-MCDM (S)
#Robots = 4	SR	265.4(212.9)	1868.0(160.1)	567.0(66.6)	737.3(61.9)
	MRv1	517.3(164.6)	1785.2(232.5)	615.9(24.2)	836.4(94.4)
	MRv2	380.3(67.7)	1780.7(233.8)	633.4(37.7)	809.3(104.2)
#Robots = 6	SR	634.4(209.3)	1978.6(200.0)	638.1(74.4)	888.0(72.6)
	MRv1	574.2(73.8)	2151.7(228.8)	702.3(52.5)	1018.3(93.3)
	MRv2	545.2(129.4)	2105.3(241.4)	701.7(18.8)	983.9(76.7)
#Robots = 8	SR	768.6(190.1)	2050.0(189.3)	708.8(56.4)	953.3(97.4)
	MRv1	606.5(207.5)	2304.5(161.6)	755.9(44.0)	1149.5(95.9)
	MRv2	540.7(139.5)	2336.9(214.5)	751.8(56.8)	1046.5(80.4)

Table 6 Results (average and standard deviation) of explored relevant area (m^2) for the mall environment, after 20 minutes of exploration. B indicates victims most likely are in big rooms, S in small rooms.

Mall		Exploration			
	Coord.	D-MCDM (B)	S-MCDM (B)	D-MCDM (S)	S-MCDM (S)
#Robots = 4	SR	2480.5(463.5)	5229.9(114.4)	2480.5(463.5)	2605.9(251.1)
	MRv1	2668.0(156.7)	4501.1(541.3)	2668.0(156.7)	2802.9(427.3)
	MRv2	2542.8(54.3)	4765.6(546.2)	2542.8(54.3)	2797.7(530.4)
#Robots = 6	SR	3164.2(478.1)	5222.3(508.8)	3164.2(478.1)	3355.1(451.2)
	MRv1	2990.1(262.1)	5449.9(349.4)	2990.1(262.1)	3706.4(436.7)
	MRv2	2941.1(198.6)	5214.6(581.0)	2941.1(198.6)	3536.6(363.2)
#Robots = 8	SR	3756.9(425.0)	5625.6(243.2)	3756.9(425.0)	3956.6(465.4)
	MRv1	3446.9(406.9)	5846.5(235.7)	3446.9(406.9)	4255.4(406.0)
	MRv2	3326.0(324.7)	5765.0(269.8)	3326.0(324.7)	3903.4(379.6)

Table 7 Results (average and standard deviation) of total explored area (m^2) for the mall environment, after 20 minutes of exploration. B indicates victims most likely are in big rooms, S in small rooms.

For example, the difference between the relevant area mapped at 20 minutes with S-MCDM and D-MCDM, in the case of victims in big rooms, with SR and 6 robots, is statistically significant (p -value = $2.42 \cdot 10^{-7}$). Fig. 6 illustrates the evolution of the explored relevant area over 20 minutes in the setting just discussed. We can observe that at the beginning the trend is almost the same for both exploration strategies. This could be explained by the fact that the 6 robots start from positions that are close to some big rooms and so also D-MCDM chooses candidate locations in big rooms. After 10 minutes, S-MCDM outperforms D-MCDM, indicating that, when there are more candidate locations in different rooms that could be selected by the robots, the benefits of using a semantic-based exploration are more evident.

Note that, similar trends are also valid for the hypothesis of victims in small rooms and, also in this case, the difference between the relevant area mapped at 20 minutes with the two exploration strategies is statistically significant (p -value = $1.34 \cdot 10^{-5}$).

For both exploration strategies, the coordination methods MRv1 and MRv2 that exploit semantic informa-




Fig. 6 Explored relevant area (m^2) over 20 minutes, in office environment, by 6 robots, with SR coordination method, in the case of victims in big rooms.

tion appear to perform relatively better than the state-of-the-art coordination method SR, and differences are statistically significant (for instance, for MRv2 vs. SR p -value = $9.24 \cdot 10^{-10}$ with S-MCDM, considering the hypothesis of victims in big rooms and 6 robots). Fig. 7 shows the explored relevant area considering the latter setting over 20 minutes. We can observe that MRv1 and MRv2 have similar trends and that they perform better than SR. This can be explained by the fact that, although there can be some initial drawbacks in sending more robots to the same candidate location, due to sensing overlaps, in the long term, there seems to be a benefit.




Fig. 7 Explored relevant area (m^2) over 20 minutes, in office environment, by 6 robots, with S-MCDM, in the case of victims in big rooms.

Only considering 4 robots, in the case of victims in small rooms, SR seems to have better results than MRv1 and MRv2, even if not statistically significant (e.g., in this setting, with S-MCDM, for SR vs. MRv2, p -value = 0.80). This similar performance of SR and MRv1/MRv2 can be explained noting that, when the number of robots is small, the exploration becomes unbalanced if more robots are assigned to the same candidate location.

Another consideration from Table 4 is that, as expected, increasing the number of robots, the amount of explored relevant area increases (apart from one degenerate case with SR and S-MCDM considering victims in small rooms and increasing robots from 4 to 6), even if the increase is not statistically significant. Note that the standard deviation of the results in Table 4 is high in the case of victims in small rooms. This could be due to the fact that, since robots should focus on small rooms, the space in which robots can move is small and, so, errors in the movements of the robots have greater influence in these experiments. Indeed, we observed in the experiments that, for example, robots can spend some time to enter in a small room.

Table 5 shows the *total* amount of explored area (as opposite to the amount of relevant area considered so far) for the office environment. The total amount of explored area increases from D-MCDM to S-MCDM in the case of victims in big rooms. For example, with 6 robots and SR, the total amount of explored area changes from 3115.6 (367.0) m^2 to 3958.0 (187.9) m^2 , with a statistically significant difference (p -value = $5.91 \cdot 10^{-6}$). Fig. 8 shows the trend over 20 minutes of such setting. This performance increase could be due to the fact that robots are encouraged to explore big rooms, from where it is possible to easily explore large portions of the environment.




Fig. 8 Explored total area (m^2) over 20 minutes, in office environment, by 6 robots, with SR coordination method, in the case of victims in big rooms.

In the case of victims in small rooms the total amount of explored area is more or less the same for D-MCDM and S-MCDM. The total amount of explored area is similar for all coordination methods. Note that the distance traveled by the robots does not change much over all the experiments (see, for example, Fig. 9). This fact shows that the difference in the amount of (relevant or total) explored area does not depend on the fact that the robots may be stuck, but almost exclusively

on the exploration strategy and the coordination methods adopted.




Fig. 9 Sum of the traveled distances (m) over 20 minutes, in office environment, by 8 robots, considering S-MCDM, in the case of victims in big rooms.

The difference in the performance of the exploration strategies can be further analyzed by looking at how they evaluate candidate locations in different rooms. As explained in Section 3.2, this evaluation for our proposed exploration strategy S-MCDM changes according to the semantic labels of the cells and to the hypothesis on the victims locations (criterion $S()$), while the state-of-the-art exploration strategy D-MCDM evaluates candidate locations in different rooms more uniformly. Fig. 10 illustrates this behavior in the case of 6 robots, SR coordination method, and victims most likely located in big rooms. This different evaluation of the candidate locations determines the number of assigned candidate locations in different rooms for D-MCDM and S-MCDM. Including semantic information in the exploration strategy effectively allows the robots to focus on candidate locations in the relevant areas, neglecting those in the irrelevant ones. Fig. 11(a) shows that, in the case of 6 robots, SR coordination method, and victims most likely located in big rooms, the number of candidate locations in big rooms assigned to the robots using S-MCDM is greater than the one in the case of D-MCDM. Fig. 11(b) illustrates that, in the same last setting, almost no candidate locations in small rooms are assigned to the robots in the case of S-MCDM.

Tables 6 and 7 show experimental results for the mall environment and report the explored relevant area and explored total area, respectively. All the above observations hold also in this environment. The only difference is relative to the case of the state-of-the-art exploration strategy D-MCDM and victims in big rooms, for which the relevant and total explored areas obtained by our proposed coordination methods MRv1 and MRv2 worsen with respect to those obtained by the coordina-

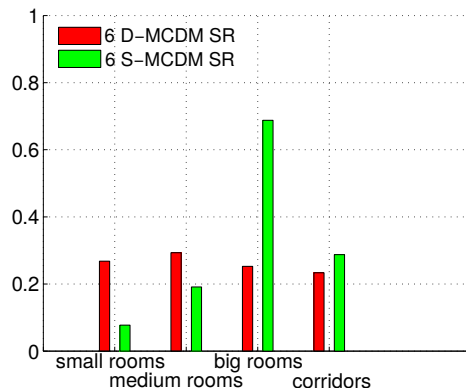


Fig. 10 Evaluation of the candidate locations (on a relative scale, average over all the candidate locations evaluated by the robots over 20 minutes) that are located in small, medium, big rooms and corridors, in office environment, with 6 robots, considering SR and the hypothesis of victims in big rooms.





Fig. 11 The number of assigned candidate locations in big rooms (a) and in small rooms (b) over 20 minutes, in office environment, to 6 robots, considering SR and the hypothesis of victims in big rooms.

tion method from the literature SR, and only with 8 robots the difference between SR and MRv2 is statistically significant (p -value=0.01). This could imply that the joint use of a coordination method that uses semantic information and an exploration strategy that does not can be inefficient.

4.4 Robustness

We experimentally verified that our results are still valid varying starting locations and the number of the robots (10 or 12). For example, Fig. 12 shows that increasing the number of robots, the explored relevant area increases. As shown in the figure, the trends for the different combinations of exploration strategy/coordination method are rather similar to those we already discussed.




Fig. 12 The explored relevant area (m²) over 10 minutes, in office environment, by 10 (a) and 12 (b) robots, with the hypothesis of victims in big rooms.

We now relax the assumption of perfect semantic information, as our system strongly relies on it. Specifically, we consider two imperfect semantic mapping modules, which make errors in assigning labels to rooms (and to cells within rooms):

- randomly according to an error rate (0.1 or 0.2 of the number of classifications), as in (Stachniss et al, 2008);
- depending on the percentage of the area actually discovered. If a candidate location p is located in a room, whose fraction of already explored area is less than a pre-defined threshold (0.2 or 0.4), the semantic mapping module classifies p randomly (with uniform probability) over the available semantic labels.

Otherwise, the semantic mapping module correctly classifies p .

We tested the system with randomly assigned semantic labels in the office environment and with victims located in big rooms. Fig. 13 shows the amount of relevant area explored over 20 minutes by 6 robots, with the random semantic mapping module. The explored relevant area diminishes compared to the case of a perfect semantic mapping module. However, the combination of our proposed exploration strategy S-MCDM and coordination method MRv1 allows to have a better performance compared to the state-of-the-art combination of exploration strategy D-MCDM and coordination method SR (at the end of 20 minutes, this difference between S-MCDM + MRv1 and D-MCDM + SR is statistically significant with p -value = $1.02 \cdot 10^{-4}$). Comparing trends of the results obtained by using MRv1, we can observe that the performance degrades, when the error rate increases. This can be explained by the fact that our proposed coordination method assigns more robots to a candidate location in a big room or a corridor, but, with an imperfect oracle, the risk is to assign more robots to areas that could be explored by only one robot. Note that in Fig. 13(a), after about 15 minutes, the performance when error rate is 0.2 becomes slightly better than that when error rate is 0.1 and this could be due to the randomness in the errors, although the difference is not statistically significant (e.g., looking at the performance at the end of the exploration: p -value = 0.2718).

Fig. 14(a) shows the amount of relevant area explored over 20 minutes, with the more realistic semantic mapping module that assigns a random label to a room if it is known less than a threshold. The performance does not degrade very much with respect to the performance obtained by our system with perfect semantic information, and S-MCDM still performs better than D-MCDM. For example, at 20 minutes, with S-MCDM and realistic semantic mapping with threshold 0.4, the explored relevant area is 1321.8 (310.2) m², while with D-MCDM and perfect semantic information, the explored relevant area is 1024.6 (220.7) m² (p -value = 0.02). The same trend is observed considering coordination methods (see Fig. 14(b)). The combination of S-MCDM and MRv1 with threshold 0.4 is still better than the state-of-the-art combination of D-MCDM and SR with perfect semantic information (1629.9 (120.8) m² vs. 1024.6 (220.7) m², p -value = $5.0 \cdot 10^{-7}$).

Finally, we tested the performance of our system by setting as termination criterion a given percentage of relevant area to be mapped (instead of the 20 minutes timeout), as in (Wurm et al, 2008). In this case,

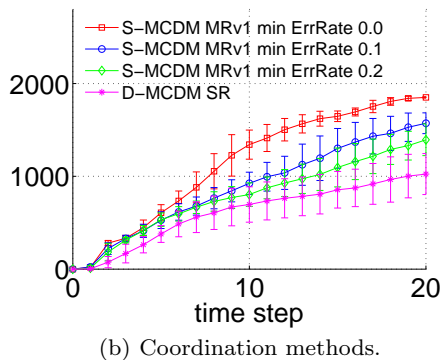
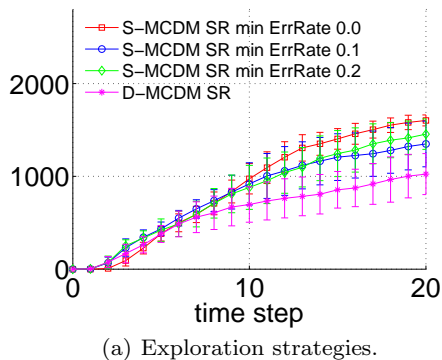


Fig. 13 Explored relevant area (m^2) over 20 minutes, in office environment, by 6 robots with random semantic mapping.


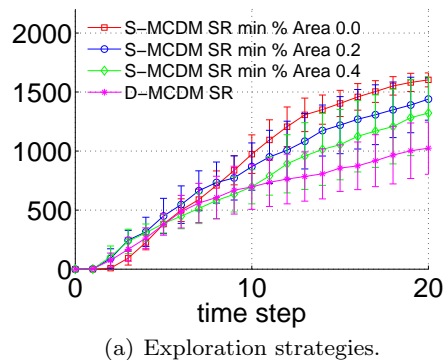


Fig. 14 Explored relevant area (m^2) over 20 minutes, in office environment, by 6 robots with realistic semantic mapping.

the system performance could be evaluated according to the time spent for accomplishing the mission. This experiment was carried out on a portion of the mall environment, with 8 robots with the goal of mapping 90% of the relevant area (victims located in big rooms). Fig. 15 shows that our proposed semantic-based exploration system with S-MCDM and MRv1 terminates earlier (around 20 minutes) than the state-of-the-art combination of D-MCDM and SR (around 29 minutes).

4.5 Discussion

In summary, results show that our semantically-informed exploration strategy largely outperforms a state-of-the-art exploration strategy in discovering areas of interest in the office and the mall environments. This can be explained by the fact that the exploration strategies that do not consider semantic information evaluate candidate locations only according to their metric features, independently of their interest for the possible presence of victims. Another relevant result is that both MRv1 and MRv2 coordination methods, which use semantic information to determine the number of robots to send to a candidate location, have better performance compared to the state-of-the-art coordination method SR.




Fig. 15 Explored relevant area (m^2), in mall environment, by 8 robots with a different termination criterion (90% of the relevant area).

This behavior is more evident with the hypothesis of victims in big rooms, because MRv1 and MRv2 directly accelerate the exploration of big rooms, as more robots are sent to such rooms. The result is valid in the hypothesis of victims in small rooms as well but, in this case, the reason seems to be that MRv1 and MRv2 send more robots in corridors, to which several rooms are connected and can be easily accessed. However, no statistically significant trend can be observed when comparing MRv1 and MRv2. In addition, our experimental results suggest that the coordination method

has comparatively less impact on the performance than the exploration strategy. This is in line with the results obtained by Amigoni et al (2012), for different search and rescue settings. Note also that our semantically-informed approach generally performs better than traditional approaches independently of the percentage of relevant area over total area. However, with few relevant areas (e.g., big rooms in office, Fig. 4(a)), the advantage in using semantic information in coordination is more evident. With many relevant areas that are easily accessible from the starting positions of the robots (e.g., small rooms in mall, Fig. 4(b)), using semantically-informed coordination is less effective (robots can be simply spread using traditional approaches with good chances of visiting relevant areas). Finally, our system proved to be enough robust to random errors in semantic labeling of the areas of the test environments.

5 Conclusions

In this paper, we have presented a semantic-based multirobot exploration approach for search and rescue that considers *a priori* information about the location of victims in order to focus on relevant areas. We have shown how to exploit knowledge of semantic map in both exploration strategy and coordination method. Experimental results obtained in two realistic test environments show that the proposed semantically-informed approach obtains significantly better performance than state-of-the-art approaches in exploring relevant areas and also, as previous work already pointed out, in exploring total area.

Future work will address the further assessment of the proposed system considering real robots with noisy communication and mapping. Furthermore, it could be interesting to change at runtime the information about relevant areas. In addition, we could find an automated way to compute some of the parameters used in our system. For example, the membership functions of the proposed coordination method can be set according to the specific building typology (e.g., being a school), on the basis of the results of Luperto et al (2013). Moreover, they could be set looking at the robots' capabilities (e.g., if sensor range is $R = 5$ m instead of $R = 20$ m, then the curves for *RoomSize* in Fig. 1 should be shifted to the left). It could be interesting also to extend this work by considering distribution of probability about the location of the victims, starting from results of Aydemir et al (2013). Moreover, a deeper study of the impact of knowledge provided by semantic maps for exploration will be performed. A direction of interest is the investigation of multi-task (MT) coordination methods (i.e., each robot plans how to reach a

sequence of candidate locations) or path optimization, starting from results of Tovar et al (2006).

References

- Amigoni F (2008) Experimental evaluation of some exploration strategies for mobile robots. In: Proceedings of the IEEE International Conference on Robotics and Automation (ICRA), pp 2818–2823
- Amigoni F, Caglioti V (2010) An information-based exploration strategy for environment mapping with mobile robots. *Robotics and Autonomous Systems* 5(58):684–699
- Amigoni F, Basilico N, Quattrini Li A (2012) How much worth is coordination of mobile robots for exploration in search and rescue? In: Proceedings of the RoboCup International Symposium, pp 106–117
- Aydemir A, Pronobis A, Göbelbecker M, Jensfelt P (2013) Active visual object search in unknown environments using uncertain semantics. *IEEE Transactions on Robotics* 29(4):986–1002
- Basilico N, Amigoni F (2011) Exploration strategies based on multi-criteria decision making for searching environments in rescue operations. *Autonomous Robots* 31(4):401–417
- Burgard W, Fox D, Moors M, Simmons R, Thrun S (2000) Collaborative multi-robot exploration. In: Proceedings of the IEEE International Conference on Robotics and Automation (ICRA), pp 476–481
- Burgard W, Moors M, Schneider F (2005) Coordinated multi-robot exploration. *IEEE Transactions on Robotics* 21(3):376–378
- Calisi D, Iocchi L, Nardi D, Randelli G, Ziparo V (2009) Improving search and rescue using contextual information. *Advanced Robotics* 23:1199–1216
- Carpin S, Lewis M, Wang J, Balakirsky S, Scrapper C (2007) USARSim: A robot simulator for research and education. In: Proceedings of the IEEE International Conference on Robotics and Automation (ICRA), pp 1400–1405
- Cipolleschi R, Giusto M, Quattrini Li A, Amigoni F (2013) Semantically-informed coordinated multi-robot exploration of relevant areas in search and rescue settings. In: Proceedings of the European Conference on Mobile Robots (ECMR), pp 216–221
- Fox D, Ko J, Konolige K, Limketkai B, Schulz D, Stewart B (2006) Distributed multirobot exploration and mapping. *Proceedings of the IEEE* 94(7):1325–1339
- Gerkey B, Mataric M (2004) A formal analysis and taxonomy of task allocation in multi-robot systems. *International Journal of Robotics Research* 23:939–954
- González-Baños H, Latombe JC (2002) Navigation strategies for exploring indoor environments. *Inter-*

- national Journal of Robotics Research 21(10-11):829–848
- Hawley J, Butler Z (2013) Hierarchical distributed task allocation for multi-robot exploration. In: Martinoli A, Mondada F, Correll N, Mermoud G, Egerstedt M, Hsieh MA, Parker LE, Støy K (eds) *Distributed Autonomous Robotic Systems*, vol 83, Springer, pp 445–458
- Howard A, Roy N (2003) The robotics data set repository (Radish). <http://radish.sourceforge.net/>
- Kuipers B, Byun YT (1981) A robot exploration and mapping strategy based on a semantic hierarchy of spatial representations. *Robotics and Autonomous Systems* 8:47–63
- Lu F, Milios E (1997) Robot pose estimation in unknown environments by matching 2d range scans. *Journal of Intelligent and Robotic Systems* 18(3):249–275
- Luperto M, Quattrini Li A, Amigoni F (2013) A system for building semantic maps of indoor environments exploiting the concept of building typology. In: *Proceedings of the RoboCup International Symposium*, pp 504–515
- Mozos OM, Stachniss C, Burgard W (2005) Supervised learning of places from range data using AdaBoost. In: *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, pp 1742–1747
- Pestman W (1998) *Mathematical Statistics: an Introduction*. de Gruyter
- Simmons R, Apfelbaum D, Burgard W, Fox D, Moors M, Thrun S, Younes H (2000) Coordination for multi-robot exploration and mapping. In: *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, pp 852–858
- Stachniss C, Mozos OM, Burgard W (2008) Efficient exploration of unknown indoor environments using a team of mobile robots. *Annals of Mathematics and Artificial Intelligence* 52(2-4):205–227
- Tadokoro S (2010) *Rescue Robotics*. Springer-Verlag
- Thrun S (2002) Robotic mapping: A survey. In: *Exploring Artificial Intelligence in the New Millennium*, Morgan Kaufmann, pp 1–35
- Tovar B, Munoz L, Murrieta-Cid R, Alencastre M, Monroy R, Hutchinson S (2006) Planning exploration strategies for simultaneous localization and mapping. *Robotics and Autonomous Systems* 54(4):314–331
- Tuna G, Gulez K, Gungor VC (2012) The effects of exploration strategies and communication models on the performance of cooperation exploration. *Ad Hoc Networks* 11(7):1931–1941
- Visser A, Slamet B (2008) Including communication success in the estimation of information gain for multi-robot exploration. In: *Proceedings of the International Symposium on Modeling and Optimization in Mobile, Ad Hoc, and Wireless Networks (WiOpt)*, pp 680–687
- Wolf D, Sukhatme G (2008) Semantic mapping using mobile robots. *IEEE Transactions on Robotics* 24(2):245–258
- Wurm K, Stachniss C, Burgard W (2008) Coordinated multi-robot exploration using a segmentation of the environment. In: *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp 1160–1165
- Yamauchi B (1998) Frontier-based exploration using multiple robots. In: *Proceedings of the International Conference on Autonomous Agents*, pp 47–53
- Zlot R, Stentz A, Dias MB, Thayer S (2002) Multi-robot exploration controlled by a market economy. In: *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, pp 3016–3023