A Particle-filter Approach for Active Perception in Networked Robot Systems

João Messias[1], José J. Acevedo[2], Jesus Capitan[3], Luis Merino[4], Rodrigo Ventura[2], and Pedro U. Lima[2]

[1] Informatics Institute,
University of Amsterdam, Netherlands, J.V.TeixeiradeSousaMessias@uva.nl
[2] Institute for Systems and Robotics,
Instituto Superior Técnico, Universidade de Lisboa, Portugal,
{jacevedo,rodrigo.ventura,pal}@isr.tecnico.ulisboa.pt
[3] University of Seville, Spain,
jcapitan@us.es
[4] University Pablo de Olavide, Spain
lmercab@upo.es

Abstract. The presence of children in a social assistive robotics context is particularly challenging for perception, mainly, in the task of locating them using inherently uncertain sensor data. This paper proposes a method for active perception with the goal of finding one target, e.g., a child wearing a RFID tag. This method is based on a particle-filter modeling a probability distribution of the position of the child. Negative measurements are used to update this probability distribution and an information-theoretic approach to determine optimal robot trajectories that maximize information gain while surveying the environment. We present preliminary results, in a real robot, to evaluate the approach.

1 Introduction

The MOnarCH project[5] focuses on social robotics in a pediatric infirmary using networked robots to interact with children, staff or visitors. This addresses explicitly the *active perception* problem as an important issue in the context of social assistive robotics for children. The issue arises when controlling a robot or group of robots so as to gather information, based on the robots sensors, that may be required by other agents (robots or medical staff). This paper considers that children carry a RFID tag and poses the active perception problem using just the RFID sensor of the robot and applied to finding a child whose location is previously unknown to the networked robot system (NRS). This is useful because robots can help to find lost or hidden children or play a hide-and-seek game with them.

The problem of active perception is that of controlling one or more mobile robots so as to maximize a given measure of information regarding a set of

[5] MOnarCH – Multi-Robot Cognitive Systems Operating in Hospitals (FP7-ICT-2011-9-601033).

features of the environment. The earliest forms of active perception focused on improving the localization of a mobile robot by controlling its motion [5]. More recently, this problem has been studied for target-tracking applications, where the uncertainty over the location of a moving target should be minimized [11]. It has also been generalized to cooperative multi-robot applications [1, 2, 4], wherein it is referred to as the *active cooperative perception* (ACP) problem.

Existing approaches to the ACP problem have formalized it in different ways under the scope of various overlapping fields of study, most notably (but not exclusively) those of Robotics, Sensor Fusion, and AI [1, 6, 11, 4]. In this family of differing approaches, the most common drawbacks are related to the scalability and generality of the proposed methods. The most accurate ACP methods depend on a careful modeling of the multi-robot system and are therefore only applicable to very specific and typically small-scale scenarios; and those that attempt to leverage larger-scale formalizations (such as those stemming from the field of AI) are forced to approximate the system and its behavior roughly, due to the complexity of the associated solution methods.

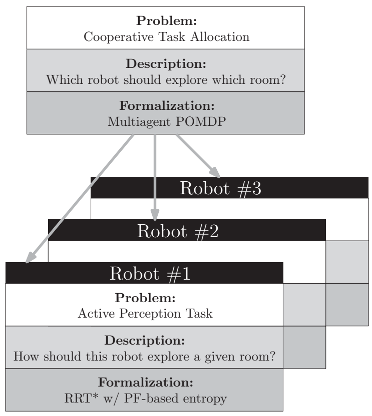

Fig. 1. A summary of the proposed ACP approach.

This paper introduces a whole ACP approach to be applied by the MOnarCH NRS, but just develops and presents results about the single robot active perception case. The ACP approach proposed in this work does not attempt to aim at optimality, but instead at being sufficiently general and easy to deploy such that it can be applied over different domains (with different environments, robots, and perceptual objectives), with minimal effort. To achieve this goal, the ACP problem is decoupled into a hierarchy of distinct subproblems (Figure 1): an allocation problem to assign a room or non-overlapping area to each robot, and a motion planning problem for each robot to decide how to explore its assigned room or area in the most efficient way.

This decomposition is also strongly motivated by the nature of the MOnarCH environment and its associated safety requirements. Specifically, in the typical use cases of the MOnarCH NRS, there should be at most one robot in a given location of interest at any time, in order to minimize the interference to the daily operations of the staff in the Instituto Portugues de Oncologia de Lisboa (IPOL).

Therefore, regarding to the proposed decoupled and hierarchic approach, at the top of this hierarchy lies a problem of multi-robot coordination: in a large environment with multiple locations of interest, what robot or sub-group of robots should be allocated to explore each of those locations? This problem has not be directly addressed in this paper, but decision-theoretic planning methods, such as Multiagent POMDPs [10], are well-suited to deal with this type of problem, although would not scale to solve the whole ACP problem for complex multi-robot domains.

At a lower level of decision-making, each robot should be able to decide how to carry out its exploration task in the most efficient way. This paper is focused on this issue, which is fundamentally a problem of motion control, which can be efficiently solved through sampling-based methods such as RRTs [3, 7]. This paper proposes a novelty approach using a Particle Filter to describe the uncertainty over the search target, and subsequently use recently proposed information-gain metrics based on the entropy of the associated non-parametric distribution. These metrics are then used to guide the probabilistic motion planning.

This document is organized as follows: in Section 2, the probabilistic planning method to generate information gathering paths is described from a theoretical standpoint, in Section 3, the preliminary results are shown, and Section 4 wraps up the paper with conclusions and a discussion on future work.

2 Active Perception via Probabilistic Planning of information-gathering paths

In this paper, we focus on the problem of planning the motion of a single mobile robot in order to search for the position of a possibly moving target.

As inputs, we are given the following:

1. The current pose of the robot;
2. A probability distribution over the target's position in the configuration space of the robot. This distribution is possibly non-parametric and typically uniform;
3. A description of the probabilistic model of the sensor of the robot that is capable of detecting the target (*e.g.* a model of the RFID reader). This description is possibly non-parametric.

The active (non-cooperative) perception problem can itself be decoupled into two interdependent problems: the problem of estimating the position of the target given the motion of the robot, and that of planning the motion of the robot in order to improve the estimation of the target's position (Figure 2).

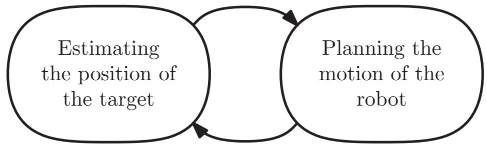

Fig. 2. The two sub-problems of Active Perception.

In the following subsections, we focus on each of these subproblems.

2.1 Estimating the Position of a Moving Target

Consider the task of estimating the position of an RFID tag carrier in an arbitrary (planar) environment, given logical-valued readings from an RFID reader mounted on a mobile robot[6]. More formally, if $p_t = [x_{p,t}\ y_{p,t}]^T$ is the position of the tag carrier at time t, $q_t = [x_{r,t}\ y_{r,t}\ \theta_t]^T$ represents the pose of the robot at that time, and $z_t \in \{False, True\}$ is the most recent RFID reading ($True$ means a positive detection and $False$ a negative one), then our objective is to minimize a given metric of confidence over the posterior probability of the target's position given the history of robot poses and measurements, *i.e.*

$$\Pr(p_t \mid q_0, z_0, q_1, z_1, \ldots, q_t, z_t)$$

This is a typical estimation / filtering problem, for which there are many applicable and well-known solutions (*e.g.* EKF, Markov localization, Occupancy-Grid methods, etc.). The "twist" in this estimation problem, as opposed to most localization and tracking applications, is that most of the readings produced by the sensor of the robot are negative w.r.t. the position of the agent. That is, if $z_t = False$, then there is a high probability that the carrier is **not** within the sensor range of the robot. Once the robot receives a positive RFID reading (or a sufficient number of positive readings to establish with a certain confidence level that the carrier is present), then the robot has succeeded in finding the carrier, and the estimation process can be terminated.

In keeping with the motivation of trying to have a general and easy-to-use solution, we have opted to approach this estimation problem using a Particle Filter (PF) to represent the uncertain position of the target.

The advantages of using a PF for this estimation problem, over its alternative methods, are the following:

- Due to the influence of negative sensor readings during the search process, the target posterior is most likely very difficult to approximate with a parametric (or kernel based) representation. The non-parametric belief representation that is characteristic of PF-based methods is ideal for this application;

[6] In this context, we are not concerned with the identity of the RFID carrier, so any identifying information contained in the RFID tag is extraneous for this problem. However, this approach could trivially be used to find **one particular** RFID tag.

– The probabilistic model of the sensor used in the search process can also be difficult to describe analytically. In the particular case of an RFID reader, the conditions for radiopropagation induced by the environment can have a complex influence on the respective readings. A PF allows for arbitrarily complex sensor models;
– The motion of the target is unknown and possibly also difficult to model in closed-form. A general PF does not make any assumptions regarding the form of this model, and random walk models could be used. In this case the standard deviation of the child movement σ is used as argument and is related to the child motion speed.

A generic PF was used for this estimation problem, following the standard algorithm which is here replicated for clarity (Figure 3). Refer to [12] for more background and details on generic particle filters.

B : Set of N particles $\langle p_x^i, p_y^i \rangle$
w : Weight vector, $w_i \in [0,1]^N$
if B or w not given **then**
  $B \leftarrow$ Sample N valid positions $s = \langle p_x, p_y \rangle$.
  $w_i \leftarrow 1/N$ for $i = 1, \ldots, N$
end if
$t \leftarrow 0$
while search not over **do**
  $B \leftarrow \mathit{predict}(B)$
  $z \leftarrow$ RFID reading at time t
  $w \leftarrow \mathit{update}(B, z)$
  if $\mathit{resample\_condition}(B, w)$ is True **then**
    $\mathit{resample}(B, w)$
  end if
  $t \leftarrow t + 1$
end while

Fig. 3. The (standard) PF algorithm used for the target position estimation.

2.2 Planning the Motion of the Robot

Given a non-parametric belief over the position of the target, which is continuously provided by the aforementioned particle filter during the execution of the ACP task, the objective of the motion planning module is to plan a path for the mobile robot over which its sensors can provide maximal information. In other words, we intend to find a sequence of poses such that the (predicted) entropy of the particle filter for target detection is minimized.

More formally, let $\mathbf{q}_{0:t} = \langle q_0, q_1, \ldots, q_t \rangle$ represent the history of poses of the robot between steps 0 and t for some $t > 0$. Analogously, let $\mathbf{z}_{0:t} = \langle z_0, z_1, \ldots, z_t \rangle$

represent the history of sensor readings. Then, our objective is to find a sequence of K *predicted* robot poses $\hat{\mathbf{q}}^*_{t+1:t+1+K}$, for some $K > 0$, such that:

$$\hat{\mathbf{q}}^*_{t+1:t+1+K} = \underset{\langle \hat{q}_{t+1},\ldots,\hat{q}_{t+K}\rangle}{\arg\min} H\left(\Pr(p_{t+1+K} \mid \mathbf{q}_{0:t}, \mathbf{z}_{0:t}, \hat{\mathbf{q}}_{t+1:t+1+K})\right) \tag{1}$$

A solution to this problem can be viewed as a form of receding-horizon control (in the field of Control Theory) or online planning (in the field of Decision Theory). We have opted to approach this problem via probabilistic sampling techniques based on the well-known Rapidly-exploring Random Tree (RRT) method [8]. This family of methods is well-suited to this problem for the following reasons:

- They are computationally efficient *anytime* algorithms (which means that after a fixed amount of time they can return the best solution found so far);
- They can handle complex kino-dynamic constraints;
- Variants of the RRT method that attempt to minimize cost functionals, such as *RRT** [7] and T-RRT, can be easily adapted to minimize the entropy of our target particle filter.

In contrast with most RRT-based methods, however, our motion planning task lacks a concretely defined goal position. Our objective is instead to find a set of poses that minimizes the entropy (1). In this sense, our approach is closest to that of [6], who have formulated information-gathering RRT variants (namely Reward-Information Gathering Trees (RIG-Tree) and Graphs (RIG-Graph)). However, the latter methods consider additional constraints over the motion planning problem (over spent energy or time) which are not necessary in our formulation. For our purposes, a simpler approach is to take a fixed-depth *RRT** or T-RRT and evaluate the cost functional (1) therein.

Although a closed-form expression for the entropy of the probability density function over the target's position is not feasible, there are suitable approximations of this measure, specifically for particle filters as in [3]. Particularly, in this work, the approximation defined in [9] based on the entropy of a Gaussian Mixture will be the used.

$$H(Pr(p_t|q_{0:t}, z_{0,t})) \leq \sum_i w_i(-log(w_i) + 0.5log(2\pi e)^2 \Sigma^4) \tag{2}$$

where w_i is the weight of the $i-$th particle of the PF, and Σ is the standard deviation for the person movement.

Note that, to calculate the entropy even with these approximations, it would be necessary to predict the state of the PF at each possible future pose. For a probabilistic sampling method such as an RRT, this would mean that a copy of the particle set of the original PF needs to be maintained at each node while expanding the search tree, so as to describe the predicted state of the PF if the robot *would* follow the path up to that node. Furthermore, that PF would need to be updated (and re-sampled) according to each possible future measurement along the path. Although this is conceptually possible, for computational reasons, we assume the following simplifications:

1. The possible future paths in the RRT always assume negative RFID measurements up to a given node. Since there is only one possible observation at each node, the tree does not need to branch out exponentially according to each possible RFID reading.
2. The motion of the robot is sufficiently fast that the motion model of the particles does not need to be considered while predicting the future state of the PF. This means that each node of the RRT do not actually have to store an instance of the full particle set (all the particles and their weights). Since, it is not required to re-sample, the positions of all the particles remain the same and each node of the RRT only need to keep their weights. By avoiding the prediction step, we also side-step the need to perform re-sampling, which is computationally expensive.

The results of our method will be discussed in Section 3.

2.3 Extension for Multiple Robots

The proposed ACP approach is such that the coordination between multiple robots would be handled at a different scope than the motion planning. However, conceptually there is nothing that precludes the extension of the above method for single-robot active perception directly to the multi-robot domain, as long as those robots can communicate freely. Since there is only one PF describing the target location, if multiple robots share the respective information, then it should be possible to use the above method to plan the paths of multiple robots simultaneously.

3 Experimental Results

The presented system has been experimentally tested in the robotics lab of the Instituto Superior Técnico in Lisbon, using the omni-directional MOnarCH robot (see Figure 4). During the experiment, the robot was in charge of finding with as much certainty as possible the location of a child carrying a RFID tag, and it had no previous information about that location. Moreover, the robot moved at 0.8 m/s and used an RFID reader to detect the person.

The right image in Figure 4 shows the more relevant information provided by the system whenever the robot requires to plan a new path to explore the environment. The PF provides the particle cloud representing the position of the person (black small dots), and the RRT* motion planner provides a connected tree. This is shown as a connected tree of green-black dots, where the green intensity corresponds to the predicted entropy at that particular pose (therefore, darker nodes are more informative).

In Figure 5, a sequence of frames taken from rviz show visually how the robot behaves during the experiment. They show the robot trajectory (blue line),

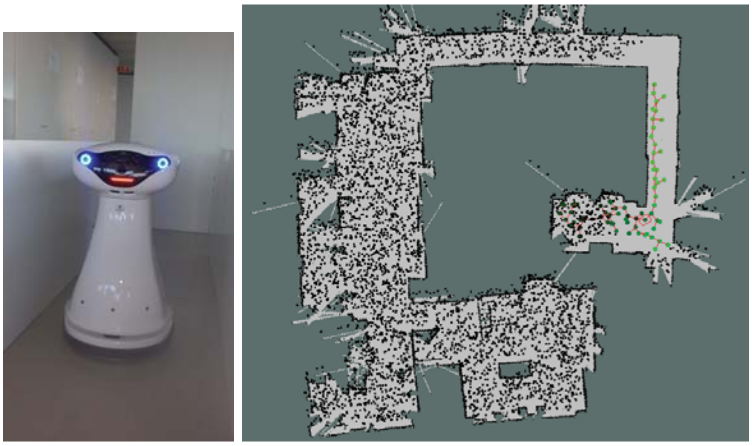

Fig. 4. Left, omni-directional MOnarCH robot used during the experiment. Right, rviz visualization of an experiment. A video from the experiment is shown in *http://youtu.be/fMB5PWQtaUI*

the last best path computed according to the proposed entropy-based RRT* (black line), the particles' distribution from the PF (black points), the particles' centroid (big red point) and the actual person location (big blue point).

While the robot moves following a computed path and getting negative RFID readings, it clears the target particles from the map. Since the robot cannot cover all particles within its sensor range and the particles are moving (based on the predict method of the PF algorithm and the proposed standard deviation of the child movement σ), these clean areas could be filled out by the nearer particles and the robot could tend to visit these areas again. So, the areas will not be fully ruled out once they have been visited, it will depend on the environment shape, the non-covered particles and the children motion speed (related to σ). As the robot has not previous information about the target location, it can start moving toward a wrong area. However, once it has explored that area without detecting any RFID tag, the robot is able to generate a new path to get back to the initial position and starts exploring other areas of the map with more uncertainty. Once the robot gets a positive RFID reading, it clears the target particles from the whole map except for the region defined by its sensor model. Following motion plans are addressed to explore again and again the same region in order to pinpointing the estimated person position.

Finally, wit respect to the features shown Figure 6, lower values imply greater certainty about the person location.

4 Conclusions and future work

The preliminary experimental results validate the proposed approach for the active perception based on a particle filter and an entropy-based motion plan-

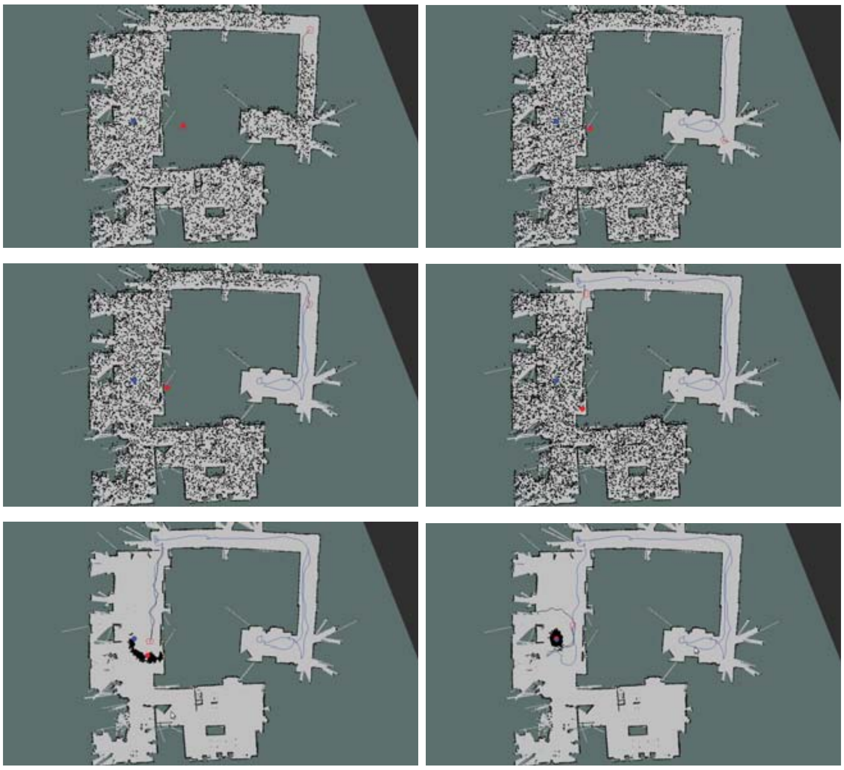

Fig. 5. Sequence of frames from the experiment rviz visualization.

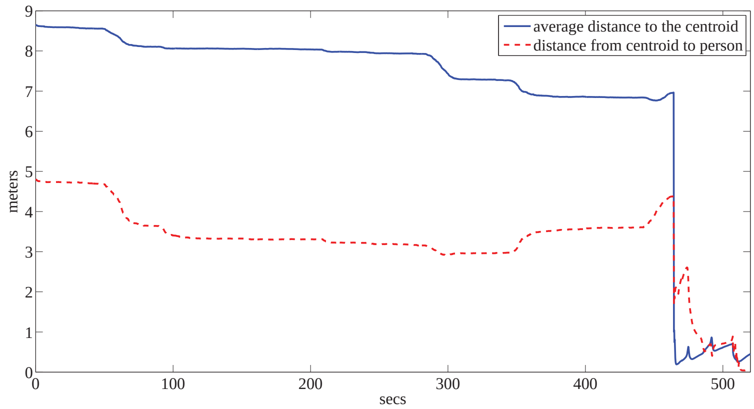

Fig. 6. This graph show the evolution during the experiment of the average weighted distance from each particle to the centroid of the particle (blue solid line) and the distance between this centroid and the actual person location (red dashed line).

ner, applied to child searching. It shows that, even having only negative measurements, the person estimator allows the system to gain certainty about the child location. Obviously, assuming complex environment with loops and children moving faster than the robots, the robots could not find the children ever. Finally, the main future developments will be directed to extend the work to multiple robots in active cooperative perception tasks, based on the dynamic area allocation and using the particle distribution.

Acknowledgments

This work is funded by the MOnarCH European project (FP7-ICT-2011-9-601033) and the Junta de Andalucia through the project PAIS-MultiRobot (TIC-7390). J.J. Acevedo, R. Ventura and P.U. Lima were also partially supported by the LARSyS (FCT [UID/EEA/50009/2013]).

References

1. Ahmad, A., Nascimento, T., Conceicao, A.G.S., Moreira, A.P., Lima, P.: Perception-driven multi-robot formation control. In: Robotics and Automation (ICRA), 2013 IEEE International Conference on. pp. 1851–1856. IEEE (2013)
2. Ahmad, A., Tipaldi, G.D., Lima, P., Burgard, W.: Cooperative robot localization and target tracking based on least squares minimization. In: Robotics and Automation (ICRA), 2013 IEEE International Conference on. pp. 5696–5701. IEEE (2013)
3. Boers, Y., Driessen, H., Bagchi, A., Mandal, P.: Particle filter based entropy. In: 13th Conference on Information Fusion (FUSION). pp. 1–8. IEEE (2010)
4. Capitan, J., Spaan, M.T.J., Merino, L., Ollero, A.: Decentralized multi-robot cooperation with auctioned pomdps. The International Journal of Robotics Research 32(6), 650–671 (2013)
5. Fox, D., Burgard, W., Thrun, S.: Active markov localization for mobile robots. Robotics and Autonomous Systems 25(3), 195–207 (1998)
6. Hollinger, G.A., Sukhatme, G.S.: Sampling-based robotic information gathering algorithms. International Journal of Robotics Research 33(9), 1271–1287 (2014)
7. Karaman, S., Frazzoli, E.: Sampling-based algorithms for optimal motion planning. The International Journal of Robotics Research 30(7), 846–894 (2011)
8. LaValle, S.M., Kuffner Jr, J.J.: Algorithmic and Computational Robotics: New Directions, chap. Rapidly-exploring random trees: Progress and prospects. CRC Press (2000)
9. Merino, L., Caballero, F., Ollero, A.: Active sensing for range-only mapping using multiple hypothesis. In: Intelligent Robots and Systems (IROS), 2010 IEEE/RSJ International Conference on. pp. 37–42 (Oct 2010)
10. de Sousa Messias, J.V.T.: Decision-Making under Uncertainty for Real Robot Teams. Ph.D. thesis, Instituto Superior Técnico (2014)
11. Spaan, M.T., Veiga, T.S., Lima, P.U.: Active cooperative perception in network robot systems using pomdps. In: Intelligent Robots and Systems (IROS), 2010 IEEE/RSJ International Conference on. pp. 4800–4805. IEEE (2010)
12. Thrun, S., Burgard, W., Fox, D.: Probabilistic robotics. MIT press (2005)