A Domestic Assistive Robot Developed Through Robot Competitions[*]

Rodrigo Ventura, Meysam Basiri, André Mateus, João Garcia,
Pedro Miraldo, Pedro Santos, Pedro U. Lima

Institute for Systems and Robotics, Instituto Superior Técnico, Universidade de Lisboa, Portugal

rodrigo.ventura@isr.tecnico.ulisboa.pt

Abstract

Scientific robot competitions are considered a powerful drive for advancing the state-of-the-art in robotics research and development. This paper describes the research objectives and technical information of the SocRob@Home team who has been representing the Institute for Systems and Robotics (ISR) from the Instituto Superior Técnico since 1998 in different robot competitions. This paper discusses in detail the hardware and software systems of our domestic robots and present algorithms used for navigation, manipulation, people and object recognition, human-robot interaction and decision making. It also presents the outcome of the participation of the team in the RoCKIn 2015 competition in Lisbon.

1 Introduction

The relevance of robot scientific competitions has been recognised in the last couple of decades as a powerful drive for pushing state-of-the-art research in robotics[Braunl, 1999; Behnke, 2006]. One of the most well-known competitions is RoboCup, encompassing both regional and global yearly events. RoboCup had its first edition in 1997 in Nagoya, initially focused on robot soccer [Kitano *et al.*, 1997]. Ever since it has significantly broaden its scope to include applications areas such as search and rescue, logistics, and domestic robots. For this latter area, the RoboCup@Home league has been focused since 2005 on the problem of mobile service robots performing tasks requested by humans in a domestic environment [van der Zant and Wisspeintner, 2007]. These competitions are based on a set of challenges of different levels of difficulty, where each team is awarded points based on their achievements. Even though the scoring is specified by a well defined set of rules, the results are not easy to replicate. Driven by the goal of defining controllable and replicable performance metrics, the RoCKIn EU project was launched in 2013 with a strong emphasis on benchmarking [Amigoni *et al.*, 2015]. Several scientific events, including two competitions, have been organised whose rules, being inspired in

RoboCup@Home, has striven to provide benchmarking both at the functionality level, e.g., navigation, perception, and at the task level, e.g., welcoming visitors, caring for comfort of an elderly person.

The SocRob team has been representing ISR/IST since 1998 in RoboCup, as the application side of SocRob (Soccer Robots or Society of Robots) ISR/IST research project. The project has involved more than 50 students over these 18 years, from early MSc to PhD students, and has reached a maturity level that enables behaviour development supported by a realistic simulator, with a GUI, where the actual code running in the robots is tested and then ported to the real hardware. Until 2013, the team's participation has encompassed Simulation, 4-Legged, Middle Size and Robot Rescue Leagues in several editions of the RoboCup World Championship and various regional RoboCup events, e.g., the Portuguese, German and Dutch Opens. Since 2013 the team decided to focus exclusively on the @Home competitions aiming towards developing service and assistive robots for future personal domestic applications.

The goal of this paper is to present the research objectives and the technical description of the SocRob@Home team in addressing the challenges posed by @Home competitions in general. The description is biased towards the RoCKIn@Home competitions, although we claim the approach is equally valid for the RoboCup@Home.

The rest of the paper is organised as follows. In Section 2 we describe our research objectives and the goals we envisage through our participation in the @Home-type robot competitions. Section 3 provides a detailed description of the robotic platform we use in the @Home competitions and Section 5 concludes the paper.

2 Research Objectives and Goals

Domestic robotics is a rapidly growing field of research, with applications ranging from simple robots for house cleaning to much smarter companion robots intended to provide care for the elderly at home. Domestic robot systems capable of providing assistance to humans must address not only traditional robotics research topics (such as sensor-fusion, task/motion planning, navigation and manipulation), but should also posses natural human-robot interaction skills. Our goal is to develop domestic robot systems, as part of a network of heterogeneous devices, to perform different duties

while interacting seamlessly with humans. To pursue such goal, we are integrating in off-the-shelf components, with outputs of our main research interests. In the following subsections we detail our research-specific objectives motivated by participation in domestic robot competitions.

2.1 Perception and Sensor Fusion

Our research in this domain includes: vision-based robot localisation [Lima *et al.*, 2011]; object tracking [Ahmad and Lima, 2013]; simultaneous localisation and tracking (SLOT) [Ahmad *et al.*, 2013]; laser-based robot localisation [Ferreira *et al.*, 2013]; and vision-based simultaneous localisation and mapping (SLAM) [Jesus and Ventura, 2013]. Particle filter-based (PF) methods have been the focus of our research, to address most perception-related problems. Using PF's, the key issues that we have been engaged in solving include: fusion of noisy sensory information acquired by mobile robots, where the robots themselves are uncertain about their own poses [Lima *et al.*, 2011] [Ahmad and Lima, 2013]; and scalability of such fusion algorithms (w.r.t. the number of robots in the team [Ahmad *et al.*, 2013]) as well as the number of objects being tracked.

For a domestic service robot working in a @Home-type environment, localisation, mapping and object/person tracking constitute the basic requirements. In addition to this, static sensors along with mobile robots in a Networked Robot System (NRS) introduce further challenging issues for sensor-fusion algorithms. Considering these, we intend to actively drive-forward our perception-related research in SocRob@Home.

2.2 Decision-Making

In prior work, we have addressed the problem of decision making for teams of autonomous robots through approaches based on the theory of Discrete Event Systems [Neto *et al.*, 2004; Neto, 2010; Costelha and Lima, 2007] and decision-theoretic formalisms for multiagent systems (Partially Observable Markov Decision Processes) [Messias *et al.*, 2011]. Recently, we have bridged these two modelling approaches, through the development and application of event-driven decision-theoretic frameworks [Messias *et al.*, 2013a; Messias *et al.*, 2013b]. The fundamental insight of this line of research is that decision making in physical environments is typically an asynchronous, event-driven process over several levels of abstraction, based on limited or uncertain sensorial information over each level, and subject to uncertain outcomes. We have explored this approach in the CMU-Portugal MultiAgent Surveillance Systems (MAIS+S) project[1], where we have successfully implemented an NRS for autonomous surveillance, comprising a team of mobile robots and a set of stationary cameras. We are currently applying some of these concepts to symbiotic interaction with autistic children and staff in a hospital, under a new CMU-Portugal project INSIDE[2].

We seek to continue our work in this topic in SocRob@Home, noting that the ability to perform decision-making under uncertainty is a fundamental requirement of any potential domestic robot, for example: given multiple tasks, such a robot must be able to manage their priorities; establish a plan for each of them; and still be able to react reliably to external events. The (possibly symbiotic) interaction with humans can also be modeled as a partially observable decision making problem. We are investigating approaches of this kind where humans interact with the robot through gestures and speech.

2.3 Human-Robot Interaction

We have focused on serviced robots in office environments, addressing in particular symbiotic autonomy: robots execute tasks requested by the users while autonomously being aware of their own limitations and asking the help of humans for overcoming them [Veloso *et al.*, 2012]. More recently, we have been moving towards speech-based communication, in order to address the @Home requirement of natural human-robot interaction. However, all communicative acts accessible from voice are also accessible through the robot touch-screen.

3 Robot Description (Hardware and Software)

Our robot builds upon a 4-wheeled omni-directional robot platform, shown in Fig. 1. This robot has been specifically developed for an ongoing European FP7 project: MOnarCH[3] [Sequeira *et al.*, 2013]. In addition to various other sensors and actuators described in [Messias *et al.*, 2014], it is equipped with two laser range finders, a Kinect RBG-D camera and a display with touch screen. On top of this platform, we installed additional devices, namely: a 7 DoF arm for manipulation (Robai Cyton Gamma 1500), a directional microphone for speech interaction (Røde VideoMic Pro), and an additional RGB-D camera (Asus Xtion PRO Live) for object detection, recognition, and localisation. The software architecture is based on ROS for middleware, while using off-the-shelf components whenever possible. This allows the team to focus on our research interests.

3.1 Navigation

Navigation is based on three modules: (1) self-localization, (2) motion planning, and (3) guidance. Our motion planning is based on a potential field approach. Instead of explicitly generating a path to the goal, it yields a potential field.

Self-localisation uses the off-the-shelf ROS package AMCL[4], employing a particle filter algorithm to fuse odometry with laser scan data. AMCL uses an occupancy grid map, which we obtain during setup time using another off-the-shelf ROS package: GMapping[5]. This package is an implementation of FastSLAM that generates an occupancy grid map from laser scan data and odometry.

Our motion planning is based on the Fast Marching Method (FMM) [Sethian, 1999] approach. Given a map constraining the workspace of the robot and a feasible goal point,

microphone

RGB-D cameras
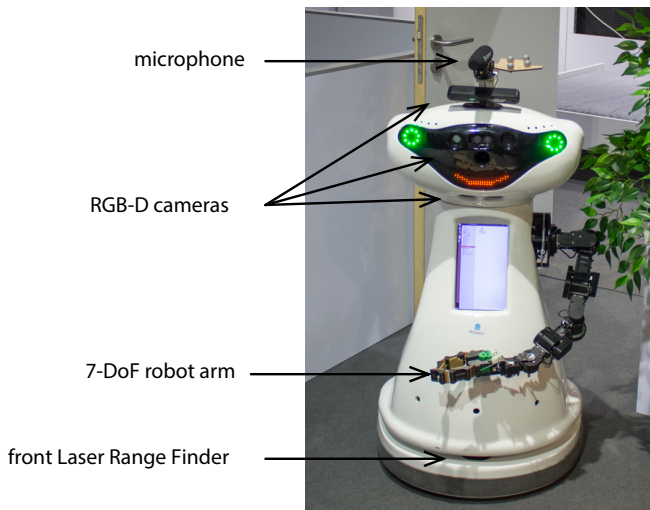
7-DoF robot arm

front Laser Range Finder

Figure 1: Robot platform of SocRob@Home.

a potential field $u(x)$, for $x \in R^2$, is constructed such that the path towards the goal is obtained by solving the ordinary differential equation $\dot{x}(t) = -\nabla u(x)$. In other words, given an arbitrary current location of the robot x, the robot should follow a gradient descent of the field $u(x)$. The use of FMM provides: local minima free path to goal across the gradient; allows the specification of a spatial cost function introducing a soft clearance to the environment obstacles; and does not require an explicit path planning and tracking. Since FMM employs a grid discretization of space, it can be directly applied to the occupancy grid map, where domain Ω corresponds to the free space in the map. Fig. 2 illustrates the results of this approach in an experiment. The cost function for the given map, allowing a certain clearance from the mapped obstacles, is shown in (a), from which, given a goal location, a field $u(x)$, is obtained (the goal corresponds to the minimum value of the field), shown in (b), and the real path taken by the robot is shown in (c).

The goal of guidance is to compute in real time the robot actuation, in terms of motion velocity, given a FMM field $u(x)$, embedding the optimal path to the goal. This is solved based on a Dynamic Window Approach (DWA) [Fox *et al.*, 1997; Brock and Khatib, 1999]. That is, given the robot's current velocity, pose, and available sensor data, the next motion velocity command is computed. This is achieved by formulating a constrained optimization problem, over a discrete set of candidate velocity commands. The outline of the algorithm is the following:

1. Generate a set of candidate linear velocity commands;
2. Discard the velocity values beyond a specified maximum absolute value;
3. Discard the velocity values which could lead to a collision (the robot is unable to stop in time, before hitting an obstacle, with the maximum deceleration);
4. Compute a fitness value for each candidate, by weighting three contributions: progress towards the goal; clearance from obstacles; and absolute speed;

5. Select candidate, maximizing the evaluation value;
6. Compute angular velocity based on the direction of the selected linear velocity, such that the robot front tends to be aligned with the motion direction.

This algorithm follows closely the DWA as initially proposed in [Fox *et al.*, 1997], except for novel methods for both computing the clearance, taking into consideration the robot shape, and the progress, based on the potential field obtained from FMM. Further details can be found in [Messias *et al.*, 2014; Ventura and Ahmad, 2015].

3.2 Manipulation

We are using Robai Cyton Gamma 1500, a 7-DoF manipulator, mounted on the base platform. The arm weight is about 2Kg, with a payload of 1500g. The drivers for ROS were re-written by the team, building on top of the low-level drivers provided by the manufacturer. Motion planning is performed by the MoveIt! library, also available for ROS. This library supports collision avoidance of the arm with obstacles (namely the robot body) during motion execution.

3.3 Interaction with users

Our platform supports two interaction modalities: touch interface over a Graphical User Interface (GUI), and speech synthesis and recognition. Text-To-Speech (TTS) employs the eSpeak[6] package while Automatic Speech Recogniton (ASR) is based on VoCon Hybrid[7], a state-of-the-art commercial solution. ASR is grammar based with the grammars created based on prior lexical knowledge of the scenarios that the robot needs to understand. To improve recognition rate, a confidence-based threshold is defined so that utterances may be discarded. Speech understanding is based on the definition of a grammar over a corpus, which spawns the possible sentences the ASR recognises.

Speech interface is currently task-oriented, i.e., the dialogue with the user is tailored towards the execution of a specific task. Our multi-modal dialogue system is based on a FSM, that coordinates the emission of canned sentences to both the TTS and the GUI, where the transitions depend on the user response (either through the ASR or the GUI). All user responses are explicitly confirmed by the robot. The outcome of each dialogue session is fed in to the main FSM, to guide the robot behaviour accordingly. An example dialogue session, taken from the participation in Rockin@Home 2015 Competition, is shown below[8]:

Robot: "Hi, would you like me to do something? I'm listening."

Human: "Close the blinds."

Robot: "I will close the blinds, is this correct?"

Human: "Yes."

Robot: "Ok, done. I'm listening."

Human: "Set the light of the living room to half."

---

[6]http://espeak.sourceforge.net/

[7]http://www.nuance.com/for-business/speech-recognition-solutions/vocon-hybrid/index.htm

[8]http://www.rockinrobotchallenge.eu/TBM3.mp4, from 0:27 to 1:05.

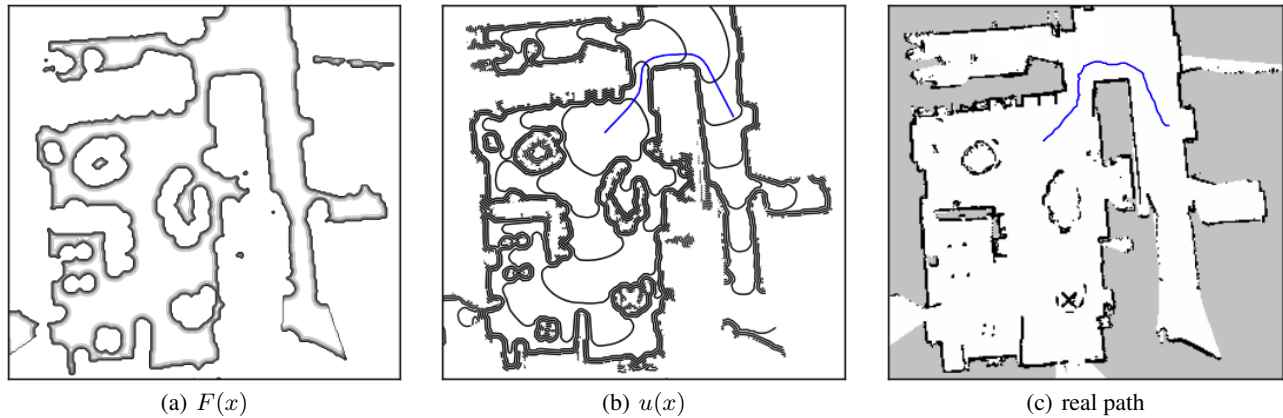|		
(a) $F(x)$	(b) $u(x)$	(c) real path

Figure 2: Motion planning using FMM: (a) the cost function $F(x)$ (darker means a higher cost), (b) the solution field $u(x)$ (level curves) together with the gradient descent $\dot{x}(t) = -\nabla u(x)$ solution (from the right to the left), and (c) the real path traveled by the robot.

Robot: "I will set the lights of the living room to half. Is this correct?"

Human: "Yes."

Robot: "Ok, done. I'm listening."

3.4   Task representation and execution

In terms of task representation and execution our approach is based on SMACH[9] which allows to represent and execute hierarchical and concurrent state machines. Our approach consists in a predefined or on-line composition of the tasks using high level state machines that address the decision-making problem and the tasks goals. High level state machines are implemented through low level state machines, with well-known outputs, based on states that address atomic functionalities like actions or conditions verification.

3.5   Object recognition

The object recognition module is based on the 3D recognition framework of the Point Cloud Library[10]. To acquire the point clouds, we use a RGB-D camera on top of the robot. It comprises two modules: a training module; and a recognition module. The training module imports models for an object class in binary PLY format. These models are then rotated and converted to point clouds, from different views, and several keypoints are extracted. The recognition process comprises three main steps: loading of the information required by the module; making the scene segmentation; and identifying clusters of objects.

The loading stage will load all the models available to the recogniser as well as specific information needed for the segmentation and coordinate conversions. This involves receiving several user defined parameters and some coordinate transformations. Next step is the segmentation of a scene's point cloud. In this step, the module will have to use either

the tabletop segmentation, for when objects are on a flat surface, or the 3D background subtraction, used (for example) when objects are on a a bookshelf. In either case the module will select the area of interest in the scene and apply a clustering algorithm to the point cloud to extract the position of the object. To classify clusters, a recogniser is trained with the previously processed models that presents the most likely correspondences.

3.6   People detection

For person recognition, we developed a ROS action server which executes two different goals. The first consists in learning someone's face, by taking multiple pictures of the face with different orientations and facial expressions. The second is the face recognition using a pre-defined number of frames that is based on the OpenCV Facerecogniser [Bradski and others, 2000]. Apart from recognizing people, the robot can also track people using the onboard MS Kinect. The algorithm consists of two steps: person identification; and tracking. The former starts by segmenting the depth image with Watershed threshold. The segmented image is then filtered to identify blobs with similar dimension to humans. The validation of the person candidates is verified with the RGB image, by training a SVM, using the Histogram Oriented Gradients, both with positive and negative samples. In order to prevent a high percentage of false positives, this method is combined with the OpenCV Haar Cascade for detecting people's upper bodies. The 3D location and speed of each blob is computed with a Kalman filter. After compensating the element's motion with the robot's egomotion (provided by the IMU), an iterative process is executed to find the best match. It compares the distances, size differences, and estimation errors among the blob considered and all the blobs in the previous frame. If the minimum value is lower than a certain threshold, it is understood that both elements (the ones in the previous and actual image) are the same, and therefore their characteristics are associated and updated.

---

[9]http://wiki.ros.org/smach
[10]http://pointclouds.org

4 Results

The participation in the camps and competitions promoted by the EU RoCKIn project[11] provided a crucial boost to leverage our competence for @Home competitions. We were in the RoCKIn Camp 2014, organized in Rome, Italy, where we received the award for "Best in Class for Manipulation"[12]. We then participated in the 'FreeBots' league in the Portuguese Robotics Open 2014, where we successfully demonstrated our robot assisting its owner in a real domestic environment[13]. In March 2015 we traveled to Peccioli, Italy, for the RoCKIn Field Exercise 2015, where our robot was demonstrated in the intelligent home of ECHORD++ RIF at U. Pisa's Service Robotics and Ambient Assisted Living Lab, winning the RoCKIn@Home Benchmarking Award. More recently, we competed in the @Home league of the RoboCup GermanOpen 2015 event in Magdeburg, where we received the "Most Appealing Robot Award".

SocRob@Home had a very successful participation in the RoCKIn 2015 final competition event[14], being the the overall @Home winners ex-aequo with the team Homer from University of Koblenz, Germany (RoboCup@Home 2015 winners), and receiving several awards in the RoCKIn@Home challenge, namely the first place among 9 participating teams in two of the three Task Benchmarks.

The RoCKIn@Home challenge is based on a user's story where an elderly person, named "Granny Annie," lives in an ordinary apartment, and suffers from typical problems of aging people. A networked robot system installed on Granny Annie's apartment helps on her daily activities. The RoCKIn@Home test bed[15] (see Figure 3) consists of a domestic environment including home automated devices (lamps, motorised blinds, and IP cameras) and equipment for benchmarking (motion capture system). Two of the task benchmarks under this umbrella story, where SocRob@Home got the first place in RoCKIn@Home 2015 are *Welcoming visitors* and *Catering for Granny Annie's comfort*. In the *Welcoming visitors* task, Granny Annie stays in bed and the robot is expected to handle different visitors appropriately, who arrive, ring the door bell, and are remotely recognised by the robot through an IP camera. In the *Catering for Granny Annie's comfort* task, the robot is called by Granny Annie using a tablet, and is expected to perform several tasks in a sequence, including remotely operating home automation devices, as well as bringing objects located anywhere in the apartment.

Besides the task benchmarks, the RoCKIn@Home challenge also includes several functionality benchmarks, targeted to quantitatively evaluate funcional capabilities of the robots: object perception, navigation, and speech understanding. As an example, we present here the results of SocRob@Home in the navigation and in the speech understanding functionality benchmarks.

In this navigation benchmark the robot is started at a predefined initial position and a sequence of 5 random reachable poses are sent through the network, one after another as they are reached by the robot. This defines 5 segments. For each segment, the travel time and the final position and orientation errors, according to the motion capture system, are recorded. Table 1 shows the results after one run of the benchmark, while the traveled path on the navigation map is shown in Figure 4. These source of these results is benchmarking data and datasets made available at the RoCKIn wiki[16]. All of these datasets follow a standardised format for the whole competition: ROS bags with specified topic names and data types.

The speech understanding benchmark comprises two stages. On the first, the teams will be given a set of spoken commands as audio files (WAV), while in the second a person will pronounce a set of commands through a loudspeaker, equally distant from all the robots. Table 2 shows the results for the two runs of the benchmark that we participated. Speech understanding is benchmarked against a Command Frame Representation (CFR) comprising a command and zero or more arguments, e.g., "MOTION(goal:"to the dining room")." The performance metrics of the benchmark are the following: FCR is the percentage of commands that the ASR engine correctly recognized, *i.e.,* the percentage of correct transcriptions; AC represents the percentage of sentences correctly parsed into the CRF format, overlooking the exactness of the transcription; WER is the Word Error Rate on the transcription of the user utterances, defined as the ratio of the number of substitutions, deletions and insertions in the transcriptions with the number of units in the sentences; and SR Acc. is the percentage of CFRs generated during the benchmark that match the gold standard CFRs (the gold standard CFR corresponds to the exact transcription and representation of the command).

More details regarding these past participations can be obtained from our team's homepage[17]. Current participation of SocRob@Home in robot competitions acts as a case study for a nationally funded research project on domestic robots[18] and for part of the research developed under the EU-FP7 project MOnarCH, which our group is currently coordinating and where most of our team members are also involved directly.

5 Conclusions

This paper described the research work and technical information of the SocRob@Home team and provided detail information about the domestic assistive robot that was developed throughout the robot competitions.

One of the most stable modules is navigation. It was developed from scratch by our team and is being used on several other projects. The module currently demanding more development effort is manipulation. The grasping effectiveness is very dependent on the camera-arm calibration and is

---

[11] http://rockinrobotchallenge.eu

[12] http://youtu.be/0STWX9SHo1I

[13] https://youtu.be/4mF0_5MCgpw

[14] Most of the footage of the video on https://youtu.be/ooO0LTiJsUc was taken from our participation in this competition.

[15] http://welcome.isr.tecnico.ulisboa.pt/isrobonet

[16] http://thewiki.rockinrobotchallenge.eu/index.php?title=Datasets

[17] http://socrob.isr.tecnico.ulisboa.pt

[18] Human-aware service robots for domestic environments (PTDC/EEI-SII/4698/2014).

(a)

(b)

Figure 3: RoCKIn@Home test bed: (a) concept; (b) real, including the trusses for the motion capture system used to obtain the ground-truth.

segment	time (s)	length (m)	avr. speed (m/s)	position error (m)	orientation error (rad)
1	16.4	3.98	0.243	0.120	0.153
2	18.4	4.86	0.264	0.213	0.152
3	29.7	10.5	0.354	0.0782	0.108
4	27.7	7.99	0.288	0.284	0.221
5	35.2	12.5	0.355	0.283	0.238
			0.301(0.0515)	0.196(0.0939)	0.1744(0.0538)

Table 1: Results from a run of the navigation functionality benchmark. The last line contains descriptive statistics of the respective columns in the format *average(std.dev.).*

Figure 4: Path traveled by the SocRob@Home robot during a run of the navigation functionality benchmark. The squares represent the initial position while the triangles denote the given waypoints.

run	FCR (%)	AC (%)	WER	SR Acc.
2	13	51	0.68	5
3	13	52	0.66	7

Table 2: Results from the two runs of the speech understanding functionality benchmark where our team participated. See text for an explanation of the performance metrics.

performed in open-loop. We are currently exploring visual servoing techniques to overcome these limitations by closing the perception-actuation loop.

We plan to use the data sets of our participation in robot competitions to benchmark our results on the major scientific challenges of the project. Nevertheless, we will recur to off-the-shelf software components, when available, to speed up the development process of our robot system, as well as to make it more dependable and competitive. These components are integrated with novel outputs of our own research using ROS.

References

[Ahmad and Lima, 2013] Aamir Ahmad and Pedro Lima. Multi-robot cooperative spherical-object tracking in 3d space based on particle filters. *Robotics and Autonomous Systems*, 61(10):1084 – 1093, 2013.

[Ahmad *et al.*, 2013] A. Ahmad, G. D. Tipaldi, P. Lima, and W. Burgard. Cooperative robot localization and target

tracking based on least square minimization. In *Robotics and Automation, Proceedings of the 2013 IEEE International Conference on*, May 2013.

[Amigoni *et al.*, 2015] F. Amigoni, E. Bastianelli, J. Berghofer, A. Bonarini, G. Fontana, N. Hochgeschwender, L. Iocchi, G. Kraetzschmar, P. Lima, M. Matteucci, P. Miraldo, D. Nardi, and V. Schiaffonati. Competitions for benchmarking: Task and functionality scoring complete performance assessment. *IEEE Robotics Automation Magazine*, 22(3):53–61, 2015.

[Behnke, 2006] Sven Behnke. Robot competitions — ideal benchmarks for robotics research. In *Proceedings of IROS-2006 Workshop on Benchmarks in Robotics Research*, 2006.

[Bradski and others, 2000] Gary Bradski et al. The opencv library. *Doctor Dobbs Journal*, 25(11):120–126, 2000.

[Braunl, 1999] Thomas Braunl. Research relevance of mobile robot competitions. *IEEE Robotics & Automation Magazine*, 6(4):32–37, 1999.

[Brock and Khatib, 1999] Oliver Brock and Oussama Khatib. High-speed navigation using the global dynamic window approach. In *Proceedings of the IEEE International Conference onRobotics and Automation*, volume 1, pages 341–346. IEEE, 1999.

[Costelha and Lima, 2007] Hugo Costelha and Pedro Lima. Modelling, analysis and execution of robotic tasks using Petri Nets. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 1449–1454. IEEE, 2007.

[Ferreira *et al.*, 2013] João Ferreira, Alberto Vale, and Rodrigo Ventura. Vehicle localization system using offboard range sensor networ. In *Proceedings of IFAC Synposium on Intelligent Autonomous Vehicles (IAV-13)*, 2013.

[Fox *et al.*, 1997] Dieter Fox, Wolfram Burgard, and Sebastian Thrun. The dynamic window approach to collision avoidance. *IEEE Robotics & Automation Magazine*, 4(1):23–33, 1997.

[Jesus and Ventura, 2013] Filipe Jesus and Rodrigo Ventura. Simultaneous localization and mapping for tracked wheel robots combining monocular and stereo vision. *Journal of Automation, Mobile Robotics & Intelligent Systems*, 6(1):21–27, 2013.

[Kitano *et al.*, 1997] Hiroaki Kitano, Minoru Asada, Yasuo Kuniyoshi, Itsuki Noda, and Eiichi Osawa. RoboCup: The robot world cup initiative. In *Proceedings of the First International Conference on Autonomous Agents (Agents'97)*, pages 340–347, New York, 1997. ACM Press.

[Lima *et al.*, 2011] Pedro U. Lima, Pedro Santos, Ricardo Oliveira, Aamir Ahmad, and Joao Santos. Cooperative localization based on visually shared objects. *RoboCup 2010: Robot Soccer World Cup XIV. Lecture Notes In Artificial Intelligence*, pages 350–361, 2011.

[Messias *et al.*, 2011] J. Messias, M. T. J. Spaan, and P. U. Lima. Efficient offline communication policies for factored multiagent POMDPs. In *Proceedings of the 25th Annual Conference on Neural Information Processing Systems (NIPS '11)*, pages 1917–1925, 2011.

[Messias *et al.*, 2013a] J. Messias, M. T. J. Spaan, and P. U. Lima. GSMDPs for multi-robot sequential decision-making. In *Proceedings of the 27th AAAI Conference on Artificial Intelligence*, 2013.

[Messias *et al.*, 2013b] J. Messias, M. T. J. Spaan, and P. U. Lima. Multiagent POMDPs with asynchronous execution. In *Proceedings of the 12th International Conference on Autonomous Agents and Multiagent Systems (AAMAS-13)*, pages 1449–1454, 2013.

[Messias *et al.*, 2014] João Messias, Rodrigo Ventura, Pedro Lima, João Sequeira, Paulo Alvito, Carlos Marques, and Paulo Carri co. A robotic platform for edutainment activities in a pediatric hospital. In *Proceedings of the IEEE International Conference on Autonomous Robot Systems and Competitions*, 2014.

[Neto *et al.*, 2004] Gonçalo Neto, Hugo Costelha, and Pedro Lima. Topological navigation in configuration space applied to soccer robots. In *RoboCup 2003: Robot Soccer World Cup VII*, pages 551–558. Springer, 2004.

[Neto, 2010] G. Neto. *Planning, learning and control under uncertainty based on discrete event systems and reinforcement learning*. PhD thesis, Instituto Superior Técnico, Universidade Ténica de Lisboa, 2010.

[Sequeira *et al.*, 2013] João Sequeira, Pedro Lima, Alessandro Saffiotti, Victor Gonzalez-Pacheco, and Miguel Angel Salichs. MOnarCH: Multi-robot cognitive systems operating in hospitals. In *ICRA 2013 Workshop on Many Robot Systems*, 2013.

[Sethian, 1999] James A Sethian. Fast marching methods. *SIAM review*, 41(2):199–235, 1999.

[van der Zant and Wisspeintner, 2007] Tijn van der Zant and Thomas Wisspeintner. *Robotic Soccer*, chapter RoboCup@Home: Creating and Benchmarking Tomorrows Service Robot Applications, pages 521–528. Itech Education and Publishing, 2007.

[Veloso *et al.*, 2012] Manuela Veloso, Joydeep Biswas, Brian Coltin, Stephanie Rosenthal, Susana Brandao, Tekin Mericli, and Rodrigo Ventura. Symbiotic-autonomous service robots for user-requested tasks in a multi-floor building. In *Workshop on Cognitive Assistive Systems, IEEE/RSJ International Conference on Intelligent Robots and Systems*, Vila Moura, Portugal, 2012.

[Ventura and Ahmad, 2015] Rodrigo Ventura and Aamir Ahmad. *RoboCup 2014: Robot World Cup XVIII*, chapter Towards Optimal Robot Navigation in Domestic Spaces, pages 318–331. LNAI 8992. Springer, 2015.