Two Faces of Human–Robot Interaction: Field and Service Robots

Rodrigo Ventura

Abstract This chapter reports on recent research and development done at Institute for Systems and Robotics (ISR) on Human-Robot Interaction (HRI), encompassing two different perspectives. On the one hand, we have conducted research on field robots tele-operated by remote human operators, targeting Urban Search And Rescue scenarios. Here, research is focused on the challenge of conveying the operator with an effective situation awareness of the robot surroundings. Situation awareness is here understood as "the perception of the elements in the environment within a volume of time and space, the comprehension of their meaning, and the projection of their status in the near future" (Drury et al. 2003). Our approach is based on decoupling operation of the robot perception from the motion command: an Head-Mounted Display displays in stereo a video stream from the cameras onboard the robot, while the gaze of these cameras is controlled by the operator head, and the robot motion is commanded from a gamepad held by the operator. On the other hand, we have addressed the challenge of service robots designed to perform tasks requested by users. Humans interact with these robots, firstly to request a task, and secondly during task execution. The main research challenge is to have these robots to be aware of their own limitations, and ask humans for help to overcome those limitations. We are currently addressing the challenge of developing robust localization and navigation methods targeting office environments.

Keywords Human-robot interaction · Urban search and rescue · Service robots · Head-mounted display · Immersive teleoperation · Interactive mapping · Fast-marching methods

R. Ventura (✉)
Institute for Systems and Robotics, Instituto Superior Técnico, Lisbon, Portugal
e-mail: rodrigo.ventura@isr.ist.utl.pt

1 Introduction

Human-Robot Interaction (HRI) is a field of study focused on the reciprocal action (interaction), including communication (verbal and non-verbal), between humans and robots [see (Fong et al. 2003; Goodrich and Schultz 2007) for reviews]. This interaction can take various forms, depending on the modality (visual, speech, etc.), on the proximity between humans and robots (collocated or remote), and on many other dimensions [see (Yanco and Drury 2002; Yanco et al. 2004) for a thorough taxonomy of the design space of HRI].

In this chapter we present recent research conducted in Institute for Systems and Robotics[1] (Instituto Superior Técnico, Lisbon) on two distinct application areas: field and service robotics. The former concerns the use of mobile robots on unstructured environments. In particular, we target Urban Search And Rescue (USAR) scenarios. In this case, the interaction is primarily concentrated on the teleoperation of the robot by a remote human operator, where the operator has total control over the robot operation. The latter area encompasses mobile robots that perform services requested by humans, where we focus on office environments. HRI happens in two situations: first, the operation of requesting a task—remotely, e.g., using the web, or locally, e.g., speech commands to the robot, and second, interaction with humans during task execution, e.g., requesting a human to open a door. The rest of this chapter is structured as two major sections, Sects. 2 and 3, covering each one of these areas of research. We wrap up the chapter with some concluding remarks on the Sect. 4.

2 Field Robotics

Field robotics concerns the use of sturdy robots in unstructured environments. By unstructured environments we mean environments for which there is no a priori map and the ground is not necessarily horizontal. A major application area where these environments can be found is Urban Search And Rescue (USAR). After a major disaster (natural or human made) affecting a urban region, victims are often found trapped inside buildings partially collapsed (and close to collapse). In many cases, these buildings pose a serious threat to human teams, and therefore, the use of teleoperated robots become a much safer option. In particular, robots can be sent in cases where human teams are unable to go due to imminent danger of collapse.

Since 2003 that ISR has been involved in the design, construction, and use of tracked wheel robots for USAR scenarios: RAPOSA, which was developed from scratch in a joint project lead by a spin-off (IdMind), and RAPOSA-NG, built upon
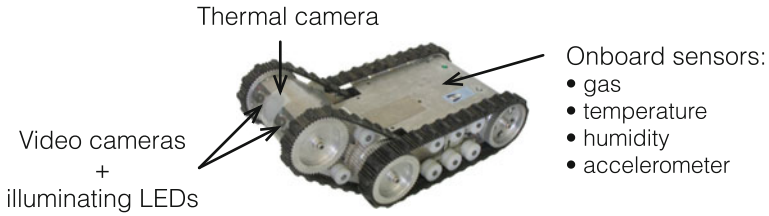
---

[1] http://www.isr.ist.utl.pt

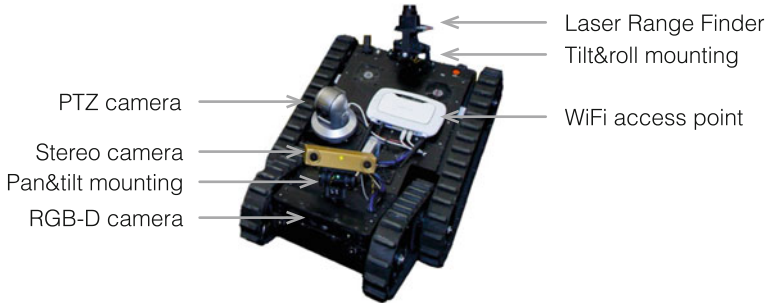**Fig. 1** Photo of the RAPOSA platform with the location of its sensors indicated

Fig. 2 Photo of RAPOSA-NG platform with the onboard equipment indicated

a commercial, barebones, and significantly improved version of RAPOSA (made by IdMind). In the following two sections these two platforms are briefly presented, followed by an overview of the conducted research on HRI issues with these platforms. Both robots have track wheel fors traction and feature two separate bodies: the main body where most of the equipment resides (batteries, motors, etc.), and an frontal arm actuated by a motor (Figs. 1 and 2). Both bodies have tracks, thus maximizing the area of traction with the gound.

2.1 RAPOSA

The first prototype was developed jointly with the IST spin-off company IdMind and the Lisbon firefighters brigade (RSBL) during the 2003–2005 period. This prototype, called RAPOSA (Fig. 1), is a tracked wheel platform whose specifications were driven by a set of requirements for USAR missions, including the capability of climbing standard-size stairs and moving along standard-size pipes, remote voice communication with victims, having environmental sensors (relevant gases, humidity, and temperature) and thermal camera. This robot is physically similar to other robots developed around the same time, notably the iRobot Packbot (Yamauchi 2004), however it is distinctive by an innovative aspect: there is a tether cable carrying power and an embedded wireless Access Point (AP), that

can be remotely attached and detached from the body without direct human intervention (Marques et al. 2007, 2006). The remote operator commands the robot to move towards the cable, using a rear camera, inside behind the cable attachment outlet of the robot. The cable is physically attached to the robot by a motorized locker mechanism that grips the cable end. The RAPOSA is equipped with a broad range of sensors: two frontal video cameras, one frontal thermal camera, one rear camera (for cable attaching), and various environmental sensors. The communication with the base station, from which the remote operator commands the robot, is based on digital wireless communications (WiFi), either directly to the base station or using the AP at the cable end.

After the development and construction of the platform, it have been intensively used for research. Besides the work on HRI, that will be described in Sect. 2.3, we have also focused on autonomous behaviors. Since USAR operations are often always conducted under a strict chain of command, the use of semi-autonomy is usually preferred over fully autonomous ones. The concept of adjustable autonomy stems from the observation that the level of autonomy is something that is decided and set by the human operator (Goodrich et al. 2001; Murphy and Rogers 1996). Under this umbrella we have addressed two autonomous behaviors: stair climbing and cable docking. We have developed a successful stair climbing method, combining visual servoing for approaching the stairs, and the use of the onboard accelerometers in closed loop to climb the stairs while keeping the robot movement orientation orthogonal with the stair edges (Ferraz and Ventura 2009). Autonomous cable docking is achieved using visual servoing techniques using the camera images captured by the rear camera (Ferreira and Ventura 2009).

2.2 RAPOSA-NG

Following the success of RAPOSA, the IdMind company developed a commercial version of RAPOSA, improving it in various ways. Notably, the rigid chassis of RAPOSA, which eventually ends up being plastically deformed by frequent shocks, was replaced by semi-flexible structure, capable of absorbing non-elastically shocks, while significantly lighter than the original RAPOSA.

ISR acquired a barebones version of this robot, called RAPOSA-NG, and equipped it with a different set of sensors, following lessons learnt from previous research with RAPOSA (see Fig. 2). In particular, it is equipped with:

- a stereo camera unit (PointGrey Bumblebee2) on a pan-and-tilt motorized mounting;
- a Laser-Range Finder (LRF) sensor on a tilt-and-roll motorized mounting;
- an pan-tilt-and-zoom (PTZ) IP camera;
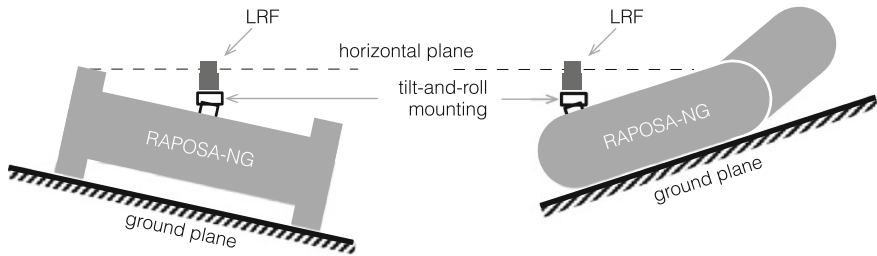- an Inertial Measurement Unit (IMU).

Fig. 3 Tilt-and-roll mounting for stabilizing the LRF

This equipment was chosen not only to fit better our research interests, but also to aim at the RoboCup Robot Rescue competition (see Sect. 2.5).

The stereo camera is primarily used jointly with an Head-Mounted Display (HMD) wear by the operator: the stereo images are displayed on the HMD, thus providing depth perception to the operator, while the stereo camera attitude is controlled by the head tracker built-in the HMD (see Fig. 4). This system will be explained in detail in Sect. 2.3.

The LRF is being used in one of the following two modes: 2D and 3D mapping. In 2D mapping we assume that the environment is made of vertical walls. However, since we cannot assume horizontal ground, we use a tilt-and-roll motorized mounting to automatically compensate for the robot attitude, such that the LRF scanning plane remains horizontal, as depicted in Fig. 3. An internal IMU measures the attitude of the robot body and controls the mounting servos such that the LRF scanning plane remains horizontal.

The IP camera is used for detail inspection: its GUI allows the operator to orient the camera towards a target area and zoom in into a small area of the environment. This is particularly relevant for remote inspection tasks in USAR. The IMU is used both to provide the remote operator with reading of the attitude of the robot, and for automatic localization and mapping of the robot.

2.3 HRI for Teleoperation

One of the challenges of using tele-operated robots in USAR scenario is to provide the remote operators with an effective situation awareness of the robot surroundings. We have been addressing this challenge by using a stereo camera pair on the robot together with an HMD equipped with an head tracker, wear by the remote operator: the stereo camera pair streams stereo imagery to the remote operator HMD, while the HMD head tracker controls the motion of the camera. With this arrangement we achieve several goals. Firstly, since the video stream is composed of stereo images, the remote operator experiences depth perception of the environment (also known as stereopsis). This allows improved perception of
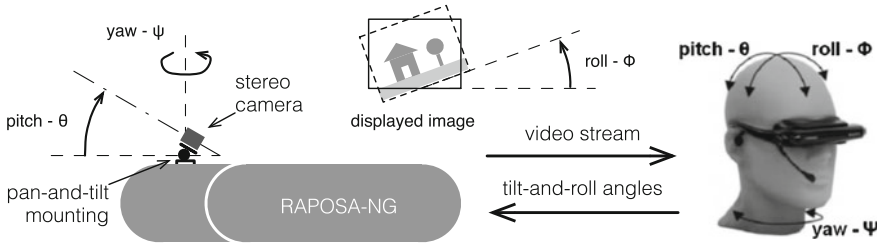
Fig. 4 Teleoperation using HMD and pan-and-tilt camera mounting

distances, than using monocular images. And secondly, by controlling the gaze of the cameras using the HMD head tracker, we allow the remote operator to control the camera orientation with his head. This frees his hands for other tasks, namely to command the robot motion using a gamepad interface.

We evaluated this architecture in both RAPOSA and RAPOSA-NG. Since the RAPOSA frontal cameras are fixed on the frontal body, we chose to orient the cameras in the following way: adjust the frontal body arm for the tilt movement, and change the heading of the robot (using the tracks) for the pan movement. We evaluated this system in a controlled user study (Martins and Ventura 2009), over a set of tasks involving RAPOSA. The results show the benefit of both stereopsis and HMD, however, they also show that controlling the heading of the robot using the HMD head tracker degrades task performance. One possible reason for this negative result lies on the fact that humans move the head and the walking direction independently. In other words, whenever the remote operator turns his head, he does not expect the moving direction of the robot to change.

Driven by this result, we decided to install the stereo camera on RAPOSA-NG using a motorized pan-and-tilt mounting. This allows the camera orientation can be changed independently than the robot heading. The HMD pitch and yaw angles control the mounting tilt and pan angles, while the roll angle rotates the images with respect to the center (see Fig. 4). Additionally, we introduced the concept of virtual pan-and-tilt: we emulate a (small) change of pan/tilt angles by re-projecting the camera image to a different optical axis (but same focal point) and cropping the result. Figure 5 illustrates the physical and reprojected optical axis. This reduces the effective field of view as seen by the operator, however, the field of view of the stereo cameras is quite large.[2] We combine the physical and the virtual pan-and-tilt so that we are able to provide fast pan/tilt image movements, up to certain limits. We evaluated this system on RAPOSA-NG in another user study (Reis and Ventura 2012), which has shown the benefits of the system.

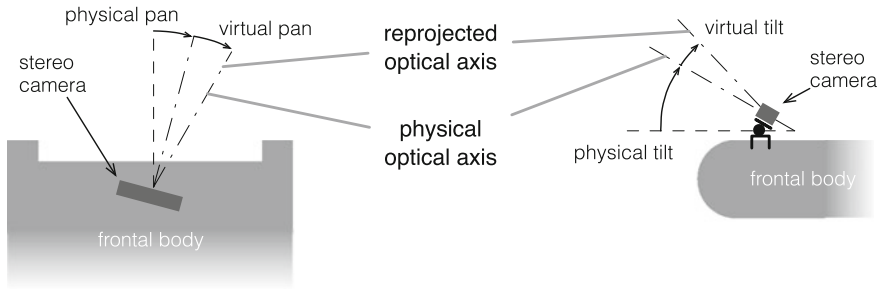[2] Pointgrey Bumblebee2 with an horizontal field-of-view of 97°.

Fig. 5 The physical and reprojected optical axis in the virtual pan-and-tilt of the stereo camera

2.4 Interactive Mapping

Most methods of 2D or 3D mapping rely on matching a sequence of sensor data scans to find the best fit among scans. Having this, sensor data is registered into a map. The most commonly used method is the Iterative Closest Point (ICP) (Besl and McKay 1992), followed by a large variety of variants. However, most scan matching methods are prone to local minima. When this happens, misalignment of scan data occurs and the map turns out showing severe distortions. We have addressed this problem taking a semi-automatic approach according to which a human user corrects these misalignments. The method is based on a graphical user interface that allows the user to interact with individual scans. Any scan can be individually translated and rotated. The innovation we have introduced lies in a technique to assist moving scans in the interface. In particular, scans can be moved along the directions of higher ambiguity (Vieira and Ventura 2012a, b, 2013). For instance, when moving scans of a straight corridor, as shown in Fig. 6, movement tends to be constrained along the corridor, while transmitting a feeling of being "stuck" on the orthogonal direction. This results from moving the selected point cloud D, with respect to M, such that a virtual force F_m caused by a mouse drag movement is balanced by a reaction force F_r given by the gradient of the ICP cost function. This effect resembles the so called "magnetic" effect in most vector drawing programs.

This methods was applied to the adjustment of both 2D scans taken from a LRF and 3D scans acquired by a RGB-D sensor. Figure 7 shows an example of a set of 2D point clouds from a segment of an office floor, (a) before alignment, (b) after running ICP, and (c) after interactive mapping. The local minima phenomena that results from ICP alone is clearly visible in the deformation of the resulting map in (b), while the involvement of a human is capable of effectively correct those deformations, as shown in (c). However, we have to acknowledge the fact that humans have prior knowledge about typical office floors, while ICP alone has not. But in fact that is the strength of having the human-in-the-loop in tasks that machines are uncapable of performing alone.
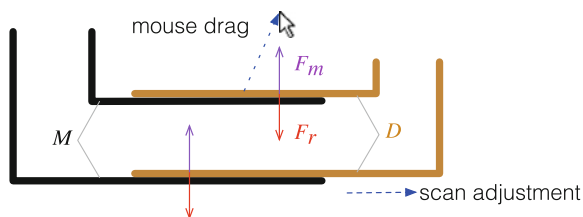
Fig. 6 Interactive adjustment of a pointcloud D with respect to M resulting from the balance between a mouse force F_m and a reaction force F_r. As a result, the pointcloud D moves along a direction parallel to the corridor
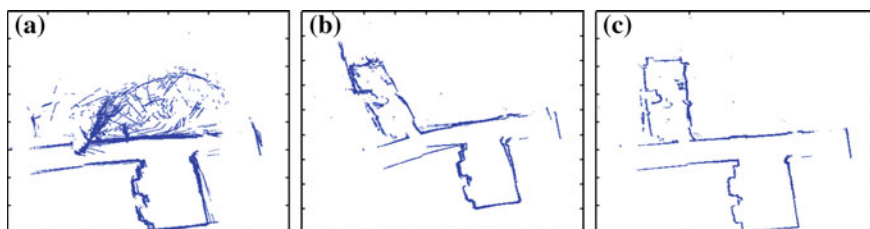

Fig. 7 Results of interactive mapping applied to 2D scans from a LRF. **a** The original raw data, **b** after running ICP, **c** after interactive mapping

2.5 Scientific Competitions as Benchmark

Evaluation and benchmarking is a fundamental part of any development cycle in science and engineering. Benchmarking aims at evaluating, with respect to an established common ground, alternative approaches to a given problem. However, many evaluation setups are *ad-hoc*, designed to specifically show in a controlled lab environment the benefits of a specific solution. In recent years, a broad range of competitions have been proposed in the area of computer science to foster benchmarking on independently controlled conditions. Examples of these competitions include semantic robot vision (Helmer et al. 2009), planning (Long and Fox 2003), and chatterbots (Mauldin 1994). In the robotics area, RoboCup has emerged as one of the most prominent scientific event comprising a broad range of competitions (Kitano et al. 1997), being robotic soccer the first one. The first competition was held in Nagoya in 1997, and since then the range of competitions has grown to include USAR—leagues Rescue Robot and Rescue Simulation (Kitano and Tadokoro 2001)—as well as service robots for home—league @Home (van der Zant and Wisspeintner 2007). All these competitions, not only allow benchmarking over a common scenario, but perhaps more importantly fosters and motivate research and development of more advanced and robust solutions. It should also be mentioned the remarkable effect of these competitions on the dissemination of science and technology to the society.

The ISR has been actively involved in RoboCup since 1998, initially in the robotic soccer leagues, and latter in the Rescue Robot one. We have participated with the RAPOSA-NG platform in the GermanOpen in 2012 (Magdeburg, Germany) and in RoboCup in 2013 (Eindhoven, Netherlands). The RoboCup Rescue Robot league evaluates the performance of USAR robots in an abridged version of the NIST standard scenario (Jacoff et al. 2003). It comprises finding simulated victims, within a limited timeframe (typically 15–20 minutes), in either a fully autonomous robot situation or a teleoperated one. When in teleoperation, the remote operator is sitting on a booth without any direct visual contact with the scenario.

These competition has pushed our development efforts along three main lines. First, in the direction of a more effective HRI, namely the stereo teleoperation with the HMD, as described in Sect. 2.2. This has proven extremely useful in competition, in particular understanding the situation of the robot in the challenging scenarios put by the NIST scenario. Second, the competition pressure demands for a real focus on the robustness of the robot. For instance, simple solutions such as remotely being able to perform an hardware reset of the onboard computer allowed us to overcome severe onboard computer hang-ups. And third, to explore the best configuration of the suite of sensors. For instance, besides the stereo cameras, we also choose to equip RAPOSA-NG with a PTZ camera for examining details that were hard to spot using the stereo camera, and with a LRF on a tilt-and-roll mounting for 2D mapping.

3 Service Robotics

Traditionally, service robots are seen as autonomous robots that perform tasks for users. The robot design is such that it is in principle capable of completing the requested task autonomously. Thus, the set of tasks is limited by the robot capabilities. For instance, if a robot has no arms, it will never be able to pick up objects by itself.

In recent years a different approach to service robots have been pursued, based on the concept of symbiotic human-robot interaction (Coradeschi and Saffiotti 2006; Rosenthal et al. 2010), according to which capability limitations of robots are coped with the help of humans. This approach has been actively pursued by the group of Manuela Veloso at CMU using the robot platform CoBot (Veloso et al. 2012). This platform is designed to navigate autonomously in office environments and to interact naturally with people. This platform is based on an omnidirectional drive base and is equipped with a suite of sensors for navigation and interaction. Research on CoBot has included a broad variety of topics, such as Kinect-based localization (Biswas and Veloso 2011), symbiotic planning (Haigh and Veloso 1998; Rosenthal et al. 2011), remote task scheduling (Ventura et al.2013), interaction with users (Rosenthal and Veloso 2010), and telepresence (Coltin et al. 2011).

Fig. 8 ISR-CoBot platform
with the onboard equipment
indicated
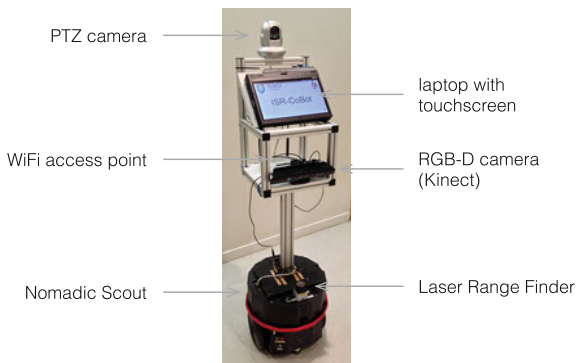
Following a collaboration with CMU on CoBot, we have started developing a
platform following the same design principles of CoBot. The platform is called
ISR-CoBot, shown in Fig. 8, being currently based on a Nomadic Scout platform
customized by IdMind. On top of it a pedestal supports the main equipment
onboard: a laptop with touchscreen for user interface, a PTZ camera for tele-
presence, a Kinect sensor, and a wireless access point.

ISR-CoBot is a differential drive platform, however we are currently working
on migrating to a omnidirectional base. The main advantage is that the kinematics
of an omnidirectional platform is the capability of deviation from obstacles
without changing the robot heading. This is a consequence of its non-holonomy
(Pin and Killough 1994).

3.1 Localization

Most localization algorithms for indoor environments match range sensor mea-
surements with a 2D map, relying on the presence of vertical walls. We have been
working on a localization method that does not depend on vertical walls, but rather
matches the map with the shape of the ground. To do so, we estimate the shape of
the visible part of the ground using a RGB-D camera, extract its edges, and match
those with the map (Vaz and Ventura 2013).

RGB-D cameras provide point clouds defined in 3D space, relative to the
camera reference frame. After an initial calibration, the geometric parameters of
the ground plane are estimated. This allows the classification of the points from a
given point cloud among two classes: ground and non-ground. Then, the concave
hull of the ground points is computed, yielding a polygon delimiting these points.
The edges of this polygon are then filtered, in order to ignore the edges corre-
sponding to the edges of the sensor field of view. The remaining ones are then
matched against a map of the environment.

3.2 Motion Planning and Guidance

For navigation we use a standard occupancy grid map (Elfes 1989), obtained from off-the-shelf SLAM software.[3] This map is used both for motion planning, using Fast Marching Method (FMM) (Sethian 1999), and localization, using off-the-shelf software.[4]

Motion planning is based on a FMM approach (Sethian 1999). Unlike other methods based on explicit path planning, e.g., RRT (LaValle and Kuffner 2001), followed by path tracking, we adopt here a potential field approach. Given a map constraining the workspace of the robot, together with a feasible goal point, a (scalar) potential field $u(x)$, for $x \in R^2$, is constructed such that, given a current robot location $x(t)$, the path towards the goal results from 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 $u(x)$. Using potential fields for motion planning was proposed in the 80's (Borenstein and Koren 1989) but they were found to be prone to local minima (Koren and Borenstein 1991). This problem can be solved by the use of harmonic potential fields (Kim and Khosla 1992), however, it does not guarantee absence of local minima at the frontier. Thus, we decided to employ a more recent approach, FMM, which provides: (1) local minima free path to goal across the gradient, (2) allows the specification of a spatial cost function, that introduces a soft clearance to the environment obstacles, and (3) does not require explicit path planning and trajectory tracking.

The FMM is based on the Level Set theory, that is, the representation of hypersurfaces as the solution of an equation $u(x) = C$. The solution of the Eikonal equation

$$\begin{aligned} |\nabla u(x)| &= F(x) \\ u(\Gamma) &= 0 \end{aligned} \tag{1}$$

where $x \in \Omega$ is a domain, Γ the initial hypersurface, and $F(x)$ is a cost function, yields a field $u(x)$ (Sethian 1999). The level sets of this field define hypersurfaces $u(x) = C$ of points that can be reached with a minimal cost of C. The path that minimizes the integral of the cost along the trajectory can be shown to correspond to the solution of $\dot{x}(t) = -\nabla u(x)$ with the initial condition of $x(0)$ set to the initial position and the initial condition $u(\Gamma) = 0$ set at the goal.[5] Intuitively it corresponds to the propagation of a wave front, starting from the initial hypersurface, and propagating with speed $1/F(x)$. FMM is a numerically efficient method to solve the Eikonal equation for a domain discretized as a grid.

---

[3] GMapping (http://wiki.ros.org/gmapping, retrieved 16 October 2013).

[4] AMCL (http://wiki.ros.org/amcl, retrieved 16 October 2013).

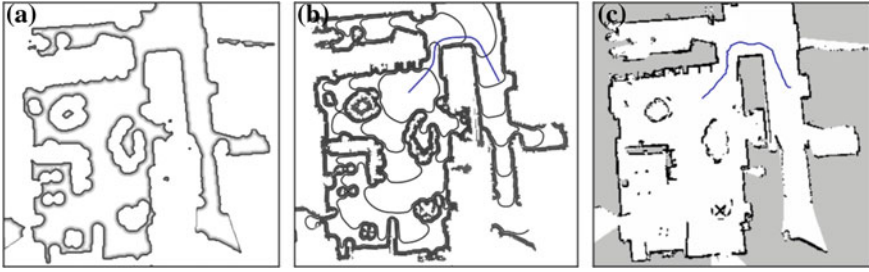[5] Γ corresponds to the boundary of an arbitrarily small ball around the goal.

Fig. 9 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

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. As cost function we use

$$F(x) = \frac{1}{\min\{D(x), D_{\max}\}} \tag{2}$$

where $D(x)$ is the distance to the nearest occupied cell in the map and $D_{\max}$ is a threshold to clip the cost function. This cost function induces a slower wave propagation near the obstacles, and thus making the optimal path to display some clearance from them. The clipping at $D_{\max}$ prevents the robot to navigate in the middle of free areas, regardless of their size. The $D(x)$ function can be directly obtained using an Euclidean Distance Transform (EDT) algorthm, taking the occupied cells as the boundary. Figure 9 illustrates the results of this approach: the cost function $F(x)$ for the given map is shown in (a), from which, given a goal location, a field $u(x)$, shown in (b) is obtained (the goal corresponds to the minimum value of the field), and in (c) the real path taken by the robot is shown.

Using FMM on a previously constructed map does not account for unmapped or moving obstacles. Thus, the field $v(x)$ used to control the robot in real-time results from combining the field $u(x)$ obtained from FMM with a repulsive potential field $r(x)$ of obstacles sensed by the LRF. This repulsive field is obtained from running EDT on a small window around the robot, such that the value of $r(x)$ corresponds to the minimum distance between any obstacle and point x. The fields are combined using

$$v(x) = u(x) + \frac{\lambda}{r(x)} \tag{3}$$

where λ is a parameters specifying the strength of the repulsive field (higher values of λ tend to increase the clearance from perceived obstacles).
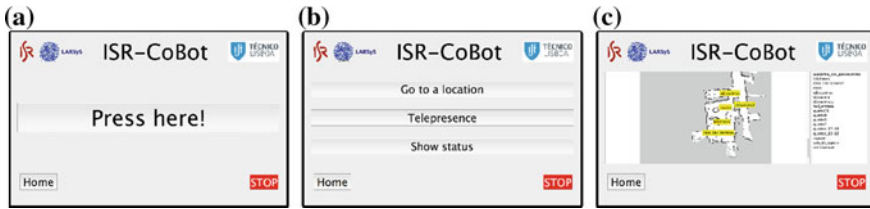
Fig. 10 Three screenshots of the onboard GUI: the welcome screen on (**a**), from which the main menu on (**b**) follows, and the destination location screen on (**c**)

The method described above have proven to be very effective, even in cluttered environments full of people crowded around the robot. We have demoed the robot in a public event—the European Researcher's Night (September 27th, 2013, in the Pavilion of Knowledge science museum, Lisbon)—where people from all ages crowded around ISR-CoBot. We have also tested the robot in an hospital environment, in the context of our participation in the FP7 European project MOnarCH (Sequeira 2013).

3.3 Interaction with Users

Currently the interaction with users is limited to setting goal locations on a displayed map of the environment, with minimal voice feedback. We decided at this stage not to use voice as input channel, since speech recognition on a mobile robot with background noise is still an open challenge (many systems are available, but they still lack the robustness needed for an engaging interaction). Therefore, even though the robot provides voice feedback of its actions in the form of canned sentences, the input is limited to the onboard touchscreen.

Screenshots of the graphical user interface can be found in Fig. 10. After a welcoming message, the user is invited to choose among three possibilities:

1. select, from either a map or a list (Fig. 10c), a location for the robot to autonomously navigate to: the shown map was previously manually labeled with relevant locations (we do not allow users to set arbitrary locations in the map, for safety reasons);
2. choose to engage in telepresence, that is, a QR-code is shown containing a URI to the Skype contact of ISR-CoBot; once called, a remote user establishes voice and visual contact with whoever is nearby the robot; and
3. display a status panel showing the current robot location and other relevant information (namely battery levels).

We demoed this interface in the aforementioned European Researcher's Night and observed that most people used appropriately the interface without any prior instruction or training.

4 Concluding Remarks

This chapter provided an overview of our current research efforts at ISR towards effective human-robot interaction. We have been exploring this long-term goal from two different perspectives. First, from the perspective of remote teleoperation of a robot, where the situation awareness by the remote user is the primary concern. In this respect, we have been using technologies and methods that aim at providing a natural interface to the operator. And second, from the perspective of autonomous service robots, which are required to robustly navigate in a structured environment, being capable of reacting appropriately to unknown obstacles, and interacting in a natural way with people.

In the future we will continue to pursue these goals. In what concerns teleoperated robots, we intend to explore augmented reality techniques to superimpose in the camera images relevant information from the map being constructed. On the service robots, we will pursue engaging interaction under the principle of symbiotic autonomy.

References

Besl, P.J., McKay, N.D.: A method for registration of 3-D shapes. IEEE Transactions on Pattern Analysis Machine Intelligence **14**(2), 239–256 (1992)

Biswas, J., Veloso, M.: Depth camera based localization and navigation for indoor mobile robots. In: RGB-D Workshop (RSS-2011) (2011)

Borenstein, J., Koren, Y.: Real-time obstacle avoidance for fast mobile robots. IEEE Transactions on Systems, Man and Cybernetics **19**(5), 1179–1187 (1989)

Coltin, B., Biswas, J., Pomerleau, D., Veloso, M.: Effective semi-autonomous telepresence. In: Proceedings of the RoboCup Symposium. Istanbul, Turkey (2011)

Coradeschi, S., Saffiotti, A.: Symbiotic robotic systems: Humans, robots, and smart environments. IEEE Intelligent Systems **21**(3), 82–84 (2006)

Drury, J., Scholtz, J., Yanco, H.: Awareness in human-robot interactions. In: IEEE International Conference on Systems, Man and Cybernetics, vol. 1, pp. 912–918 (2003)

Elfes, A.: Using occupancy grids for mobile robot perception and navigation. IEEE Computer **22**(6), 46–57 (1989)

Ferraz, J., Ventura, R.: Robust autonomous stair climbing by a tracked robot using accelerometer sensors. In: Mobile Robotics: Solutions and Challenges (CLAWAR-2009), pp. 415–422. World Scientific (2009)

Ferreira, F., Ventura, R.: Autonomous docking of a tracked wheels robot to its tether cable using a vision-based algorithm. In: Workshop on Robotics for Disaster Response, IEEE International Conference on Robotics and Automation (ICRA-09). Kobe, Japan (2009)

Fong, T., Nourbakhsh, I., Dautenhahn, K.: A survey of socially interactive robots. Robotics and autonomous systems **42**(3), 143–166 (2003)

Goodrich, M.A., Schultz, A.C.: Human-robot interaction: A survey. Foundations and Trends in Human-Computer Interaction **1**(3), 203–275 (2007)

Goodrich, M.A., Olsen, D.R., Crandall, J.W., Palmer, T.J.: Experiments in adjustable autonomy. In: Proceedings of IJCAI Workshop on Autonomy, Delegation and Control: Interacting with Intelligent Agents, pp. 1624–1629 (2001)

Haigh, K.Z., Veloso, M.M.: Interleaving planning and robot execution for asynchronous user requests. Autonomous Robots **5**(1), 79–95 (1998)

Helmer, S., Meger, D., Viswanathan, P., McCann, S., Dockrey, M., Fazli, P., Southey, T., Muja, M., Joya, M., Little, J., Lowe, D., Mackworth, A.: Semantic robot vision challenge: Current state and future directions. arXiv preprint arXiv:0908.2656 (2009)

Jacoff, A., Messina, E., Weiss, B.A., Tadokoro, S., Nakagawa, Y.: Test arenas and performance metrics for urban search and rescue robots. In: Proceedings of IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS-2003), vol. 4, pp. 3396–3403 (2003)

Kim, J., Khosla, P.: Real-time obstacle avoidance using harmonic potential functions. IEEE Transactions on Robotics and Automation **8**(3), 338–349 (1992)

Kitano, H., Tadokoro, S.: RoboCup rescue: A grand challenge for multiagent and intelligent systems. AI Magazine **22**(1), 39–52 (2001)

Kitano, H., Asada, M., Kuniyoshi, Y., Noda, I., Osawa, E.: RoboCup: The robot world cup initiative. In: Proceedings of the First International Conference on Autonomous Agents (Agents'97), pp. 340–347. ACM Press, New York (1997)

Koren, Y., Borenstein, J.: Potential field methods and their inherent limitations for mobile robot navigation. In: Proceedings of the IEEE International Conference on Robotics and Automation (ICRA-91), pp. 1398–1404 (1991)

LaValle, S., Kuffner Jr, J.: Randomized kinodynamic planning. The International Journal of Robotics Research **20**(5), 378–400 (2001)

Long, D., Fox, M.: The 3rd international planning competition: Results and analysis. Journal of Artificial Intellligence Research (JAIR) **20**, 1–59 (2003)

Marques, C., Cristovão, J., Lima, P., Frazão, J., Ribeiro, M.I., Ventura, R.: RAPOSA: Semi-autonomous robot for rescue operations. In: Proceedings of IROS2006 — IEEE/RSJ International Conference on Intelligent Robots and Systems (2006)

Marques, C., Cristovão, J., Alvito, P., Lima, P., Frazão, J., Ribeiro, M.I., Ventura, R.: A search and rescue robot with tele-operated tether docking system. Industrial Robot **34**(4), 332–338 (2007)

Martins, H., Ventura, R.: Immersive 3-d teleoperation of a search and rescue robot using a head-mounted display. In: Proceedings of 14th IEEE International Conference on Emerging Technologies and Factory Automation (ETFA-2009) (2009)

Mauldin, M.L.: Chatterbots, tinymuds, and the turing test: Entering the loebner prize competition. In: Proceedings of AAAI-94, pp. 16–21 (1994)

Murphy, R.R., Rogers, E.: Cooperative assistance for remote robot supervision. Presence **5**(2), 224–240 (1996)

Pin, F.G., Killough, S.M.: A new family of omnidirectional and holonomic wheeled platforms for mobile robots. IEEE Transactions on Robotics and Automation **10**(4), 480–489 (1994)

Reis, J.J., Ventura, R.: Immersive robot teleoperation using an hybrid virtual and real stereo camera attitude control. In: Proceedings of Portuguese Conference on Pattern Recognition (RecPad) (2012)

Rosenthal, S., Veloso, M.: Using symbiotic relationships with humans to help robots overcome limitations. In: Proceedings of the AAMAS'10 Workshop on Collaborative Human/AI Control for Interactive Experiences. Toronto, Canada (2010)

Rosenthal, S., Biswas, J., Veloso, M.: An effective personal mobile robot agent through symbiotic human-robot interaction. In: Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems, vol. 1, pp. 915–922. International Foundation for Autonomous Agents and Multiagent Systems (2010)

Rosenthal, S., Veloso, M., Dey, A.: Task behavior and interaction planning for a mobile service robot that occasionally requires help. In: Proceedings of the AAAI Workshop on Automated Action Planning for Autonomous Mobile Robots. San Francisco, CA (2011)

Sequeira, J., Lima, P., Saffiotti, A., Gonzalez-Pacheco, V., Salichs, M.A.: MOnarCH: Multi-robot cognitive systems operating in hospitals. In: ICRA 2013 Workshop on Many Robot Systems (2013)

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

van der Zant, T., Wisspeintner, T.: Robotic Soccer, chap. RoboCup@Home: Creating and Benchmarking Tomorrows Service Robot Applications, pp. 521–528. Itech Education and Publishing (2007)

Vaz, M., Ventura, R.: Ground-plane based indoor mobile robot localization using RGB-D sensor. In: Proceedings of Portuguese Conference on Pattern Recognition (RecPad) (2013)

Veloso, M., Biswas, J., Coltin, B., Rosenthal, S., Brandao, S.D., Kollar, T., Meri cli, C., Samadi, M., Ventura, R.: CoBots: Collaborative robots servicing multi-floor buildings. In: IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS-12). Vila Moura, Portugal (2012)

Ventura, R., Coltin, B., Veloso, M.: Web-based remote assistance to overcome robot perceptual limitations. In: AAAI Conference on Artificial Intelligence (AAAI-13), Workshop on Intelligent Robot Systems. AAAI, Bellevue, WA (2013)

Vieira, P., Ventura, R.: Interactive 3D scan-matching using RGB-D data. In: IEEE International Conference on Emerging Technologies and Factory Automation (ETFA'12). Kraków, Poland (2012a)

Vieira, P., Ventura, R.: Interactive mapping in 3D using RGB-D data. In: IEEE International Symposium on Safety Security and Rescue Robotics (SSRR'12). College Station, TX (2012b)

Vieira, P., Ventura, R.: Interactive mapping using range sensor data under localization uncertainty. Journal of Automation, Mobile Robotics & Intelligent Systems **6**(1), 47–53 (2013)

Yamauchi, B.M.: PackBot: a versatile platform for military robotics. Proceedings of SPIE **5422**, 228–237 (2004)

Yanco, H.A., Drury, J.L., Scholtz, J.: Beyond usability evaluation: analysis of human-robot interaction at a major robotics competition. Human-Computer Interaction **19**(1), 117–149 (2004)

Yanco, H.A., Drury, J.L.: A taxonomy for human-robot interaction. In: Proceedings of the AAAI Fall Symposium on Human-Robot Interaction, pp. 111–119 (2002)