Received December 16, 2014, accepted December 29, 2014, date of publication April 17, 2015, date of current version April 27, 2015.

Digital Object Identifier 10.1109/ACCESS.2015.2424273

Rubbing Shoulders With Mobile Service Robots

FERAS DAYOUB, TIMOTHY MORRIS, AND PETER CORKE, (Fellow, IEEE)
ARC Centre of Excellence for Robotic Vision, School of Electrical Engineering and Computer Science,
Queensland University of Technology, Brisbane, QLD 4000, Australia

Corresponding author: F. Dayoub (feras.dayoub@qut.edu.au)

This work was supported by the Australian Research Council through the Discovery Projects Scheme under Project DP110103006.

ABSTRACT This paper is not about the details of yet another robot control system, but rather the issues surrounding real-world robotic implementation. It is a fact that in order to realize a future where robots coexist with people in everyday places, we have to pass through a developmental phase that involves some risk. Putting a ‘‘Keep Out, Experiment in Progress’’ sign on the door is no longer possible, since we are now at a level of capability that requires testing over long periods of time in complex realistic environments that contain people. We all know that controlling the risk is important—a serious accident could set the field back globally—but just as important is convincing others that the risks are known and controlled. In this paper, we describe our experience going down this path and we show that mobile robotics research health and safety assessment is still unexplored territory in universities and is often ignored. We hope that this paper will make robotics research labs in universities around the world take note of these issues rather than operating under the radar to prevent any catastrophic accidents.

INDEX TERMS Health and safety, mobile robots.

I. INTRODUCTION

So far the major role for robots in our life has been realised on the factory floor. Factory robots are a 50-something year old technology. They are generally immobile and operate within human restricted zones. More recently, we are witnessing a great interest in robots for non-factory environments that are slowly but surely taking more and more roles in our daily life from cleaning the floors of our homes [1], guiding us in museums [2] and morphing into being our cars [3]. Eventually, we think — as many roboticists do — that our relationship with these robots will become more intimate to the extent that they will become personal care-givers and private companions. We want and need these robots to perform many low-skill or high-risk jobs which make human life easier and improve its quality. In short, we strongly believe that mobile autonomous robots are heading to become part of our everyday life, in the same way as phones and tablets are today, and they will cause a similar revolution in the way we live our lives. We are passionate believers in this future, but to get to this future from the present day means that testing and development of robots should at some point leave closed laboratories and head to the public realm. This raises a number of interesting considerations that robotics research groups in universities around the world will sooner or later have to deal with.

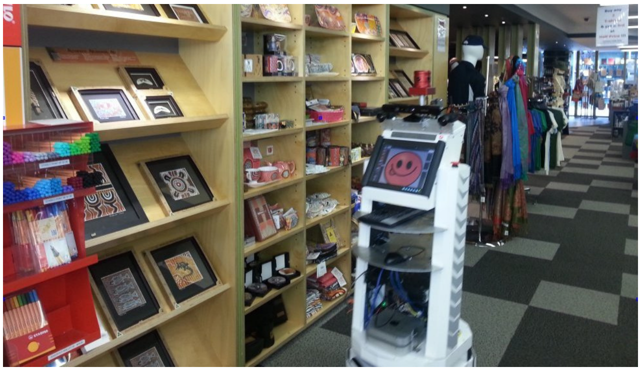

FIGURE 1. **Our robot in a bookshop.**

Our own research is concerned with lifelong autonomy [4], [5] for vision-guided robots, that is, how can robots using predominantly visual information operate for very long periods (weeks to months) in environments that are typical of the real world — complex, crowded and time varying in terms of structure as well as visual appearance. We are also committed experimentalists, and evaluating the efficacy of such a system requires very long term experimentation. Our modus operandi for experiments, like we suspect for many other robotics researchers, had been to stay back one night, run the experiment and make the ICRA video clip. Like most other researchers we operated on the well-known principal of it being better to ask forgiveness than permission. However, to achieve the desired long-term experimentation, night work and operating under the radar

TABLE 1. Summary of the risk assessment which we had to undertake before we were allowed to conduct our experiment on campus.

Hazard	Likelihood	Risk associated	Control measures
Collision with a human standing in front of the robot	Unlikely	Injury to persons involved	**a)** Speed limit the robot. The robot is equipped with a safety micro-controller. This controller serves as a watchdog for several components in the robot. Speed control, bumper contact, and laser scanner driver all emit watchdog signals verified by the safety controller. If any of these signals is not received, or if the wheel speed exceeds preset threshold, the safety controller halt the robots motion. **b)** The robot is equipped with multiple kill switches (E-Stops). The E-Stop switches can be triggered either manually or automatically by triggers installed in the environment, cutting power to the motors.
Blocking access to emergency exits	Possible	Obstructing evacuation	**a)** No-GO areas in front of emergency exits and lift doors have been installed using magnetic kill strips on the ground and reflective tape on the ceiling. **b)** In the case of a fire alarm, the operator of the robot has the responsibility of stopping the robot and moving it away from emergency exits. **c)** Inform the relevant Fire wardens about the robot.
Servicing the robot while batteries are connected	Unlikely	Injury to persons involved	**a)** Only qualified people to service the robot. Testing and tagging of the charging unit. **b)** All operators to follow safe operating procedure and pre-run checklist.
Robot falls down the stairs and falls on to somebody	Rare	Injury to persons involved	**a)** No unsupervised operation in areas with opens stairs unless redundant automatic emergency stop triggers are installed in front of the stairs doors.**b)** $1m^2$ No-GO areas in front of stairwell doors opening have been installed using ceiling and floor safety strips.
Batteries might rupture from overheating or improper use	Rare	Injury to persons involved	The batteries are concealed inside the robot and not connected directly to the charger. There is a circuit between the charger and the batteries to prevent overcharging and also to switch the robot off when the batteries charge drop below a pre-set level. The user does not have direct access to the batteries.
Person attempting to push robot	Possible	Injury to persons involved	Use signs on the robot to tell people not to try and physically stop the robot and to use the E-stops instead.The motors disengage as soon as an E-stop button is pressed after that pushing the robot is easy.
Eye Damage laser scanner	Rare	Injury to persons involved	Use certified eye safe lasers.

was just not going to cut it. The experiments needed to be very long term, operating in daylight and night-time illumination conditions (to test the robustness of computer vision) and they needed to be public, very public, since people are a key part of most environments. There was no alternative, we took a deep breath and asked permission — this is our story.

II. THE ISSUES

The issues, and considerations, are more complex than we initially envisaged. However, we can break them down into four important categories: safety of people, safety of the robot, ethics and university rules and regulations. While the details will vary across countries and institutions, the underlying principles are, we believe, widely applicable.

A. HUMAN SAFETY

As is increasingly common around the world, universities are growing more strict about occupational health and safety. Universities have always had a duty of care to staff, students and visitors. However, in recent years, in Australia at least, the legal bar has been raised which means that academics can be taken to court if negligent. Universities are of course accustomed to dealing with many risks and hazards, for example oxidising and toxic chemical, dangerous biological materials, ionising radiation and so on. However, when it comes to robots we were in unexplored territory and the safe/conservative response was just to say "don't do it".

1) HEALTH AND SAFETY ASSESSMENT

The first step was conducting a health and safety evaluation along with a risk assessment, and a summary of our assessment is presented in Table 1. In our assessment the worst outcome was the robot entering the stairwell and falling down the stairs. A falling 75 kg robot has the potential to inflict very considerable injury to any person in the stairwell. So, we modified the environment to include redundant and automatic means of emergency stopping the robot in defined spatial regions — creating *no-go zones*. These are discussed in Section III-C.

Interestingly, health and safety officials were more concerned with the robot colliding with people or running over their feet. We were urged to add flashing lights and an audible alert (continuous beeping), such as found on automated machines in some industrial environments. We argued very hard to avoid this, figuring that the flashing and beeping would make the robot very unwelcome in our work environment, and the flashing light would play havoc with its visual perception. We won that argument, but it put the onus back on us to "prove" that the robot would not collide with people. To avoid collisions with people or objects the robot has several 3D perception systems, discussed in Section III-A, as well as bump sensors on its base and sides. The maximum speed of the robot is limited by a low-level micro-controller which is programmed to clamp motor velocities at a predetermined maximum, in our case 0.4m/s, which is non threatening. As a last resort the robot has a number of emergency stop buttons. In practice it has been easy and reliable to navigate around people moving in the corridors, even when it is very busy during semester time.

A more difficult problem, surprisingly, is avoiding students who are not moving, but sitting in the corridor. Late in semester students study in the corridors, sitting on the floor, leaning against the walls, with their legs out. These are very low obstacles, below the level of the laser scanner, and we need to deal with this case, see Section III-D.

2) DEMONSTRATING SAFETY

In order for the robot to be able to operate unattended, we had to convince ourselves and the university that it was safe to do so. This was challenging since there are no health and safety precedents for long-term unattended operation of robots on campus, and several recent (non robotic) health and safety incidents exacerbated the university's concern. Ultimately the issue of operating a novel technology in an institution is one of confidence. We know the system and its levels of performance but how do we convince somebody else of its efficacy. The answer was to go through a probationary period where we employed a student during their semester break to supervise the robot and maintain a log book of events and run time. Once the probationary period was completed without any major incident we no longer employed a robot supervisor but rather left it to the researcher to periodically check on the robot throughout the day. By doing so we were able to build confidence and move to the next stage of unattended operation. Table 2 shows a typical log book entry.

TABLE 2. **An entry in the operation logbook during human watch stage.**

Run	Date	Start Time	End time		Time	Reason
5	22/8/13	8:30am	6:00pm	contacts	12:31pm	student kicked bumper

3) EMERGENCY PROCEDURES

Everybody at the university needs to comply with emergency procedures. In the event that a floor needs to be evacuated (e.g. in the case of a fire) an alarm sounds and then everybody leaves via the emergency exits. We were originally asked to have the robot detect an emergency situation, ensure it was not blocking any exit doorway and then shut-down. This seemingly simple task is one that we did not believe we could perform in a reliable and fail-safe manner. Detecting the alarm is probably feasible using sound recognition software (but we haven not investigated this) but moving away from any exit doors requires that the robot is correctly localised and the navigation software operating correctly, neither of which we could guarantee under all circumstances. Instead we use the *no-go zone* mechanism, described in Section III-C, to keep the robot out of large regions in front of emergency exits so that it can never block the exit.

The final step was to train the fire wardens whose job during an emergency is to ensure that everybody has left the floor. We trained them in how to e-stop the robot if it was being problematic and, once e-stopped, to push it out of the way.

4) EDUCATION

A final component of human safety was education. The robot has several conventional red e-stop buttons as well as bump plates on its sides and around its base. While we know these are present and are comfortable enough to use them, we could not expect anybody to else to even know this capability exists. We created a very large (A0) poster, see Figure 2, which we placed on the wall of the lift lobby of our floor so that it was the first thing that any visitor would see. We tried to keep it simple and positive, discuss safety but keep it non threatening. The key principle we tried to convey was to *"treat the robot like a person and respect its 'personal' space"*. The poster includes a description and diagram of where it has bump sensors and most importantly where the kill-switch buttons are located to activate the emergency stop. On this poster we reassure people that these buttons can be pressed should they simply feel uncomfortable with the robot. In practice there has not been a problem with people e-stopping the robot.

B. ROBOT SAFETY AND WELFARE

Our robot needs to be able to recharge itself automatically, and for its own well-being not escape via the stairs or the lift (elevator). Our biggest concern was that the robot would move into the stairwell and fall down the stairs, or get into the lift and disappear.

1) ON-BOARD SAFETY MEASURES

Our robot includes a *Watchdog* micro-controller which is programmed to clamp motor velocities at a predetermined maximum. It also detects failure of high-level software should it stop sending regular motor commands. It also monitors bump sensors which temporarily drop motor power and large red human-pressable emergency stop buttons, which cut power to motors entirely. A physical key is required to restart.

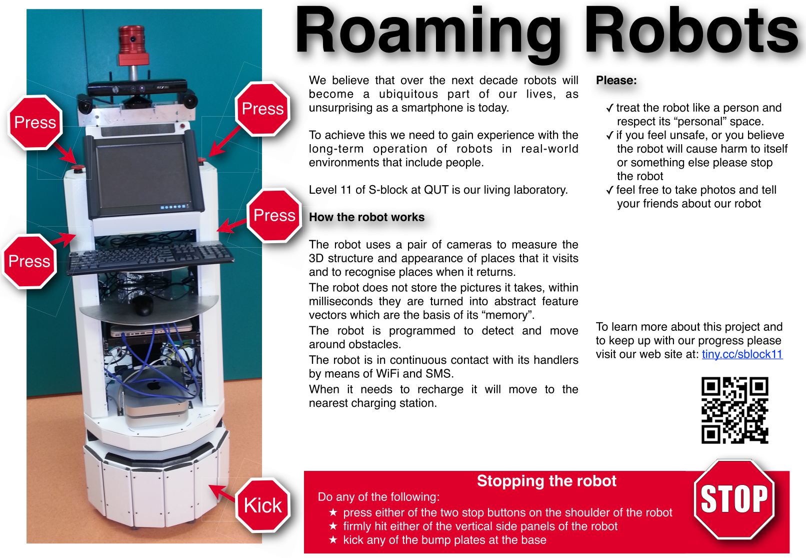

Roaming Robots

We believe that over the next decade robots will become a ubiquitous part of our lives, as unsurprising as a smartphone is today.

To achieve this we need to gain experience with the long-term operation of robots in real-world environments that include people.

Level 11 of S-block at QUT is our living laboratory.

How the robot works

The robot uses a pair of cameras to measure the 3D structure and appearance of places that it visits and to recognise places when it returns.

The robot does not store the pictures it takes, within milliseconds they are turned into abstract feature vectors which are the basis of its "memory".

The robot is programmed to detect and move around obstacles.

The robot is in continuous contact with its handlers by means of WiFi and SMS.

When it needs to recharge it will move to the nearest charging station.

Please:

- ✓ treat the robot like a person and respect its "personal" space.
- ✓ if you feel unsafe, or you believe the robot will cause harm to itself or something else please stop the robot
- ✓ feel free to take photos and tell your friends about our robot

To learn more about this project and to keep up with our progress please visit our web site at: tiny.cc/sblock11

Stopping the robot

Do any of the following:

- ★ press either of the two stop buttons on the shoulder of the robot
- ★ firmly hit either of the vertical side panels of the robot
- ★ kick any of the bump plates at the base

FIGURE 2. Poster to educate visitors about the robot. It is placed prominently and is the first thing people see when they exit the lift.

2) NO-GO ZONES

We use these zones to ensure human safety (e.g. clearing emergency exits and avoiding stairs), robot safety and security (e.g. avoiding stairs and lifts). Automatically detecting these defined spatial regions in a fail-safe manner is difficult. In addition to that, the health and safety officers put the condition that the robot should not come within 1m from stairs entrances, elevator's doors and emergency exits which prevented us from using a sonar array for drop-off detection in the case of stairs for example.

Our solution is to use redundant and independent sensors on the robot which are triggered by different modifications to the environment: magnetic strips in the floor and retro-reflective on the ceiling. We also use these zones to ensure the safety of parts of the environment, for instance in one of our test areas there are very large touch screens that need to protect. More details about can be found in [6].

3) RECHARGING

Our robot docks with a wall mounted recharging station. This station is around 300 mm heigh and requires the robot to press against it to reveal the charging contacts. To successfully activate the charging station requires a controlled forward collision of the robot with the dock. During this procedure the collision avoidance system is overridden and objects that pass between the robot and dock might be collided with. The bump switches will still remain active, halting the robot, but of course any unnecessary collision is undesirable. To handle this situation the robot plays a "fog horn" sound when initiating docking/undocking to alert people between the robot and the dock. Then, a beeping tone similar to a truck reversing alert continue until the procedure is complete. Handling undocking was simpler as the robot has rear sonar sensors that can detect rear obstacles.

4) PERVERSE HUMAN BEHAVIOUR

Everybody seemed to readily accept the robot, it rapidly became part of the environment and was largely ignored. However, we did notice one group of students who would frequently, for sport, crowd the robot and force it to drive into the wall. The sport of herding was unfamiliar to the robot as the problem of navigation is naively implemented as strict obstacle avoidance between goal points. The robot's ignorance of social conventions such as variable limits of personal space and right of way meant that it was unable to demonstrate its intentions and escape. Our open-day experiments also highlighted the need for the robot to show firmly but safely its intent, and this is an important area for future research. For the time being we can only try and educate people that even robots have personal space, see Figure 2.

C. ETHICS

As robotics researchers our focus was on the technological aspects of the problem as described above. However ethics, privacy and research data management are all important considerations.

1) INTERACTION WITH HUMANS

Our robot operates in an environment with people (academics and students) and *interacts* with them in a simple fashion. Does this level of *interaction* constitute human experimentation? Is there an ethics issue here? This was a bigger worry to us than the technology, since we are familiar with the technology, but not with ethics and the ethics approval process.

2) PHOTOGRAPHY OF HUMANS

There is a growing debate around privacy, particularly with respect to the trails we leave in our online lives. There has been specific concern about privacy and robots, caused by the ubiquity of low-cost quadrotor platforms that can carry cameras. In the University context we have to comply with the National Privacy Act [7].

> F/6.2 Information privacy:
> The Information Privacy Act applies to personal information which is defined as information or an opinion, including information or an opinion forming part of a database, whether true or not, and whether recorded in a material form or not, about an individual whose identity is apparent, or can reasonably be ascertained, from the information or opinion.
> ... Where information is recorded in a way which cannot be linked to a known individual and the personal information has become de-identified, then the privacy principles do not apply.

and the concern is regarding all the images captured by a visually-guided mobile robot operating in a public place. In order to film students for publicity purposes, for example, we need to obtain permission and a signed consent form. The images captured by the robot would very likely include people (students, staff, visitors) in the field of view, and they would be identifiable in the image, and the time and place would also be recorded. We respond to this concern in two ways. Firstly, the images are very rapidly turned into features (e.g. SURF) from which it is no longer possible to recognise people, there will be no privacy issue. We decided to be upfront about this, and explain all this on the robot's web site and on the informational poster Figure 2. Secondly, there are times when we wish to log raw images to use for subsequent algorithm development or comparison, and in this case we have to ensure that the data is saved on a secure university server, see next section. We have agreed not to publish any image with close up facial shots without the prior permission from the subjects in the image. In the absence of agreement faces would be blurred, a solution also adopted by Google Street View.

3) MANAGING RESEARCH DATA

Our university has policies regarding research data, again driven by national laws. Everything recorded by the robot is considered research data and therefore has to comply with the following policies

> *D/2.8 Management of research data:*
> *Research data means data in the form of facts, observations, images, computer program results, recordings, measurements or experiences on which an argument, theory, test or hypothesis, or another research output is based. Data may be numerical, descriptive, visual or tactile. It may be raw, cleaned or processed, and may be held in any format or media.*
> *D/2.6 QUT Code of Conduct for Research:*
> *All research data, including primary materials, are considered to be University records and must be stored, disposed of or transferred in accordance with the QUT records management policy (F/6.1).*
> *... When collecting, storing, using or disclosing personal information, researchers must abide by the mandatory requirements of the Information Privacy Act 2009 and the University's Information privacy policy (F/6.2).*

D. ORGANISATIONAL ISSUES

After addressing all the issues above, the safety of people, safety of the robot, ethics and privacy, we still needed a permission to run the robot outside our labs. At first glance this seems like asking for trouble, but there are real issues around insurance and we wanted to go for full disclosure. We don't actually *own* the spaces that we think of as our own, that is our offices, laboratories and connecting corridors. They are managed variously by schools and faculties but are ultimately all owned by central facilities management. After some difficult initial conversations we were given the permission we sought. What helped was the clear preparation that we had done, our thinking through of all the issues, and great advocacy from the school and faculty. The bookshop, see Section V-C, is leased and required different approval, but they were extremely welcoming of the robot and saw it as a real talking point.

III. TECHNOLOGICAL ASPECTS

A. THE PLATFORM

Our experimental platform is a MobileRobots' Research GuiaBot shown in Figure 3. The robot has differential steering, weighs 75 kg and can run for 1 hour after about equal charging time. The sensors include two Point Grey Grasshopper monochrome cameras which each have a 1.4M pixel 2/3″ CCD sensor and a FireWire interface, forward looking SICK laser range finder, an RGBD camera (XBOX 360 Kinect) and backward looking sonar sensors. The robot has three on-board computers all running ROS (Robot Operating System) [8].

FIGURE 3. **Left: our robot, Guiabot, navigating a hallway. Right: the charging station where the robot homes autonomously to charge itself.**

To achieve robust localization and reliable obstacle detection we built a laser-based occupancy maps using

GMapping [9]. For localization, the robot uses `AMCL`, an adaptive particle filter localizer [10] provided in ROS. The navigation algorithm consists of a global and a local planner. The global planner generates the complete path to a goal using an Dijkstra algorithm. The local planner is seeded by the global plan and generates velocity commands to control the robot. For further information see [11]. Finally, for reliable obstacle detection, the laser scans were augmented by 3D data from a Kinect sensor as described below.

B. NAVIGATION BETWEEN PEOPLE

People who are standing up and walking are easy to navigate around and we experienced relatively little problem here. We used a strategy of treating both people and structure as objects that occupy space that the robot plans around to avoid collision. If a situation occurs when no free space exists to plan to the goal the robot stops until it had free space in the direction of the goal. When blocked for sufficient time the robot will give up for now and choose a new goal.

C. NO-GO ZONES

No-go zones are installed to ensure safety during unmonitored autonomous operation. The areas of particular concern are lift entries, stairwells and fire escape doors. We defined no-go zones within 1 m of these places, see Figure 4. The robot is equipped with redundant and independent mechanisms to stop it should it enter a no-go zone. The first is an redundant array of Hall effect sensors under the base of the robot. They are activated by magnetic strips (Neato boundary marker) installed under the carpet. The second is an upward pointing industrial photoelectric sensor which closes a contact if a piece of retroreflective tape is within a specified range. These sensors were wired directly into the emergency stop system and thus independent of any user provided software. The e-stop logic will cause a stop if the wiring to the sensors is broken, and the sensors themselves signal stop if they are de-powered. Once an e-stop has occurred human intervention is required to restart the robot: after checking for the cause of the failure a physical key must be put into the switch and turned. With this system installed the robot will become disabled at a safe distance from the hazard if high-level navigation failed to avoid them.

FIGURE 4. **Redundant safety measures to stop the robot driving into no-go zones. We can clearly see the patch of different carpet under which the magnetic strips are layed, as well as the retroflective tape on the ceiling.**

D. SENSOR FUSION FOR SAFETY

Previous supervised deployments relied completely on cameras that provided a single forward facing point of view. This resulted in reaction only to objects within the cameras effective Field of View (FoV). The effective FoV for detecting objects using the camera head is a subset of the overlapping 57°/82.4° horizontal and 43°/66.9° vertical FoV provided by the Kinect and stereo camera respectively. To detect objects closer to the robot we made a slight downward tilt to the cameras resulting in reliable ground detection approximately 2 m from the base.

During supervised operation occasional intervention was required to avoid collision with objects outside the camera FoV. To reduce the chance of collision during unattended operation we use the forward facing laser as a collision sensor of last resort. The laser provides a 180° horizontal FoV mounted 0.3 m from the ground. This was fused with the information provided by the onboard cameras to improve obstacle detection (see Figure 5).

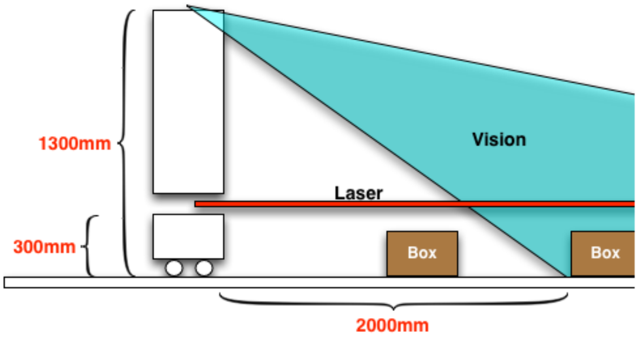

FIGURE 5. **Illustrates the sensor overlap of the vision head at 1.3 m and the laser mounted at 0.3 m.**

Collisions may occur with objects such as thin legged chairs and bookshelves with recessed panelling (Figure 6) at the height of the laser, giving a false footprint, meaning the closest parts of an object have been undetected. Fusing the two sensors therefore improves reliability but our current

FIGURE 6. **The bookshelf has a recessed gap at the laser sensor height giving the bookshelf a false footprint.**

configuration shown in Figure 5 is still unable to detect an object shorter then 0.3 m at a range less then 2 m. We would expect, in normal operation, to plan around these objects at sufficient range. The most common occurrence of these problematic low height objects were small boxes and students sitting in hallways with their legs stretched out, typically giving the robot plenty of time to react from a distance. If there is an unintended collision one or more of the bump panels surrounding the robot's base are activated, disengaging the motors temporarily, giving the robot a chance to re-evaluate its options.

In practice, even for indoor operation the Kinect sensor has interesting failure modes. One of our test environments is panelled with touch screen display and their infrared emission renders the Kinect blind (Figure 7). In conclusion, a diverse range of sensor modalities employing different physical measurement principles is critical for robustness and safety.

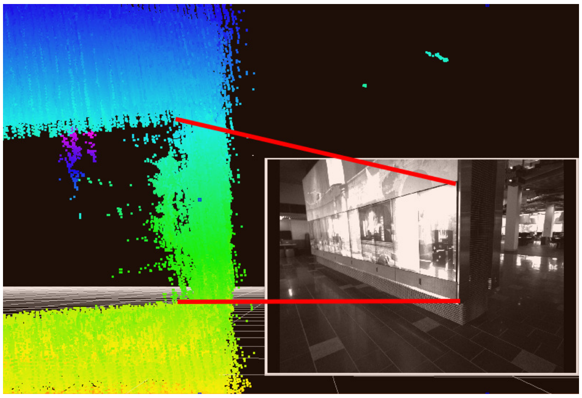

FIGURE 7. **Kinect range image of a flat wall with a touch screen display that emits infrared, the sensor returns no range.**

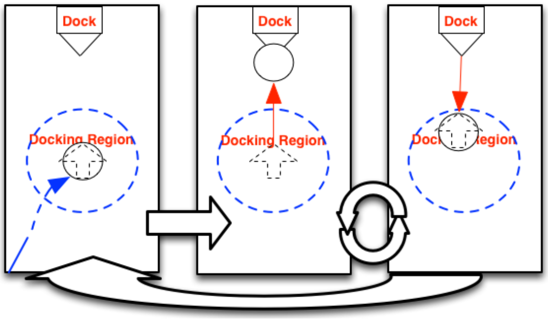

FIGURE 8. **Left: The robot navigates to a goal facing the dock approximately 1 m out. Middle: The robot servos in using a prior laser appearance of the dock. Right: If a charging connection is not made correctly the robot reverses out a maximum of 1 m to retry servoing again. If after a minute a connection has not been made the robot navigates again to the docking goal in an attempt to improve the starting location.**

E. AUTOMATIC RECHARGING

To conduct long-term experiments meant that we had to consider automatic recharging. Figure 8 illustrates the states in the docking and undocking procedure. Undocking reverses the robot approximately 1 m to a pose which is stored as the *docking goal*. When the robot needs to recharge, the current goal is replaced with the *docking goal*. On reaching the goal, the normal planning and obstacle avoidance are temporarily suspended and give way to a servoing algorithm based on the prior appearance of the dock using laser. This is why it is important to remember and return to the *docking goal* as this leaves the robot facing the dock ready for the servoing algorithm to begin a controlled collision into the dock itself. When a collision with the dock is detected by the bump sensors, an attempt is made to draw current via the onboard contacts. If the robot and the dock do not align correctly, the robot uses rear sonar to reverse without collision. Servoing into the dock is then repeated. Due to the nature of the docking procedure, which is a controlled collision with the docking station, the robot had to disable the obstacle avoidance during that period. To insure safety the robot play an alarm sound to notify people around the robot to give way if they are standing between the robot and the dock.

On occasion we found the robot attempting to dock with a corner of the neighbouring vending machine. This was due to localisation error when returning to the *docking goal* and the similar profile of the vending machine edge. The error typically arose when the robot was crowded near the docking location while people waited to use the lifts and vending machine. Therefore if the robot is unable to align its contacts and draw current from the charger after a minute of trying, it returns again to the *docking goal*. To return to the *docking goal*, goal-based planning and obstacle avoidance systems are reactivated. This navigation gives the localiser a chance to recover and improve the starting location for the servoing procedure to begin again. This procedure is repeated until successful charging is detected.

IV. MOBILE SERVICE ROBOTS IN THE REAL WORLD

Over the past years, several long-term experiments with mobile robots have been conducted inside public places.

The earlier examples are the museum tour guides RHINO [12] and MINERVA [13] which appeared in 1998. In a confined environment, these robots operated for seven days covering 19 km and two weeks covering 44 km. Soon after that the Mobot Museum Robot Series [14] have appeared where a number of robots were deployed for five months in a public museum. These robot were able to autonomously navigate for days at a time with automatic re-charging. In order to simplify the navigation task an achieve such long-term autonomy, artificial landmarks were installed in the environment. The museum tour guide robot trend continued with the RoboX robots [15]. Around ten robots were deployed during The Swiss National Exhibition'02 to guide visitors. Another example is TOOMAS the Shopping Guide Robot [16]. Nine robotic shopping guides were deployed in three different home improvement stores in Germany. Another example and in order to demonstrate a robust autonomous navigation inside an office environment, a long-term experiment using a PR2 robot was conducted by [11]. The robot operated for 30 hours covering 42 m at a 0.4 m/s. There was no automatic recharging. When charging was required, the robot parked near the charging station and asked an operator to plug it in. Very recently, the four CoBots robots [17] have jointly reached the 1000 km marks of autonomous navigation inside a university building over a period of three years.

V. OUR EXPERIENCES

Our robot has clocked up to 150 kilometres of running in three quite different environments over a period of six months. This distance was covered while the robot navigating autonomously including an autonomous docking to the charger station when low on battery. We ran the robot during typical working hours of 9am-5pm. The following section give more details about the nature of each environment.

A. OFFICE ENVIRONMENT

Our robot normally inhabits a quiet environment on a research floor of S-Block, one of the buildings in QUT Gardens Point campus. During operating hours, the robot roams the floor greeting students with its smiley face and automatically recharging when the battery level drops. Figure 3 shows the robot in S-Block during some vision-only navigation experiments [4]. People in this environment are robot savvy: academics, postgraduate students and final year students in the robotics, aerospace and mechatronic disciplines. Navigation challenges in this environment include low light levels, long sections of texture free walls and glass walls. On average, the robot covers 2.8 km per day.

This environment has many offices with manual doors, which our unarmed robot cannot open. Some doors have swipe card access which would require a robot with two arms to operate. Therefore we were limited to goal locations in the corridor only. We eventually obtained the ability to override the swipe access doors and prop them open which gave the robot access to more places and a choice of charging stations.

In this environment, the robot played the role of a receptionist with its touch screen displaying the names of different staffs, after the selection of a name, the robot guides the visitor to the desired office. When low on battery, the robot returns to one of the two charging stations and repeat autonomously.

B. ROBOTRONICA

Robotronica was the first public robotics exhibition in Australia which took place at the Queensland University of Technology on 18 August 2013 and drew around 24,000 people [6]. Figure 9 shows part of the exhibition floor where the robot operated. This floor contains over 40 multi-touch screens facing a large lounge area containing tables and chairs and glass-walled classrooms.

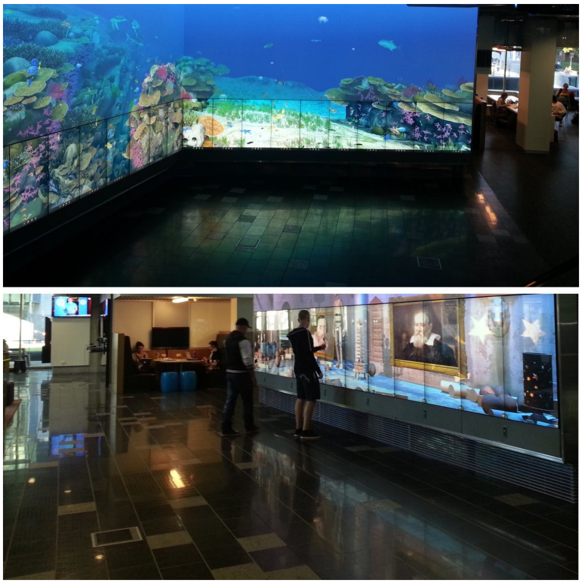

FIGURE 9. **The CUBE, one of the world's largest digital interactive and learning environments with over 40 multi-touch screens. This photo was taken in the early morning before the crowd showed up.**

While the robot deals well with glass walls in its normal environment, the large video displays confound the vision-based map building and obstacle detection since the static world assumption is violated. In addition, we encountered slotted walls that the SICK laser beam would pass through and large touch screen walls that emit infrared rendering even an RGBD camera blind in areas (Figure 7). We used a fusion of sensors to compensate for the environments challenges, although in this deployment the robot was under constant supervision and therefore did not require additional infrastructure.

The robot was given a list locations for some of the attraction on the exhibition floor and was required to cycle between them. For most of the day the robot struggled to find room to move (Figure 10). The average minimum distance to obstacles was 0.67 m and it was only able to navigate a distance of 987 m. A video of the robot in the exhibition can be found at http://youtube/ZJgDB3nu4zs. The robot had to be escorted by a human to the charging dock, it would have run out of power before it got there unaided.

The important insight taken away from this deployment was that static safety margins neglect the fact that in crowded places people are willing to share their personal space and they in turn expect that they can "rub shoulders" with the robot while it is moving.

FIGURE 10. **The Guiabot among people during the Robotronica event.**

C. THE BOOKSHOP

Most recently the robot has been roaming the campus bookshop (Figure 1) and outside cafeteria. The robot task in this environment was to greet the customers and interact with them through its touch screen. The customers can select one of different book sections in the bookshop then the robot navigate them to that section.

The location has significant dynamics and highlight the challenges of operating in the real world [5]. Our current semi-supervised experiments test mapping and navigation in this highly dynamic environment where efficient paths can change over time. The paths chosen by the robot under different modes of operation will be compared over a long period, along with assessments of crowd density made by the human supervisor. This information will be used to measure navigation performance of a robot that adapts to its environment versus one that does not.

The biggest challenge in this environment were low shelves below the plane of the SICK laser, as discussed in Section III-D. We found the robot bumping into objects placed outside the 2D plane of the laser (around 0.3 m from the ground). To address this issue, we enhanced the robot's Field of View (FoV) by augmenting 3D information provided by a forward facing camera head on top of the robot, approximately 1.3 m above the ground. This enabled objects below the laser height to be detected and avoided successfully.

VI. EMERGING GLOBAL SAFETY STANDARDS FOR SERVICE ROBOTS

In this article we have told the story about how we convinced ourselves and others of the safety of long-term unsupervised mobile robot operation in a public space. It was a lot of work and clearly shows that this is not a trivial undertaking. However, with robotics rapidly developing and new applications being addressed we are clearly entering a new era where mobile service robots have to co-exist with humans in everyday environments, which makes a global safety standards for handling such robots become a necessity.

There is an international standardization effort aimed squarely at this problem and standardising risk assessment guidelines — ISO has recently released an International Standard document 13482 [18]. The realization of this global standard safety will have a positive impact on the market of personal robots as it provides essential protection for small startup companies which was previously missing. This protection allows companies to innovate and explore new possibilities and at the same time conduct the necessary risk assessment for their product in a logical and thorough way. This is very important in the event of litigation around an accident with a new robotic product. Prior to ISO 13482 a company would need not only to perform a risk assessment, but significant legal resources to prove in a court of law that they have conducted a fit and proper risk assessment. The new standard sets the benchmark for what constitutes a proper risk assessment for a mobile robot.

A significant benefit of this new standard is the comprehensive hazard identification analysis that it provides. The standard covers 85 potential hazards and provides safety measure to control them. In spite the fact that some of the hazards mentioned are not relevant to our case, all the hazards that we have thought of are mentioned in the document. If we had this standard at the beginning of our project we would have had a much easier time communicating with the health and safety officers who are responsible of approving our risk assessment and that would have also given them more confidence in allowing us to carry out our research experiments.

VII. CONCLUSION AND DISCUSSION

In order to realise a future where robots co-exist with us in everyday places, we have to pass through a development phase that involves some risk. With a new technology controlling the risk is important, a serious accident could set the field back globally, but just as important is convincing others that the risks are known and controlled. Putting a "Keep Out, Experiment in Progress" sign on the door is no longer possible since we are now at a level of capability that requires testing over long periods of time in complex realistic environments that contain people.

In this article, we have described our experience going down this path and found that robotics health and safety assessment is still unexplored territory. In order to advance our research we needed to interact with non-practitioners and convince them that we knew what we were doing, and that it was safe. We had to go through the process of creating a risk assessment ourselves and creating the controls. We strongly believe that robotics research labs in universities around the world should take note of these principles, to do the right thing in the lab and not operate under the radar.

Perhaps the robotics research community should work together to produce a standard risk assessment procedure and a code of conduct specifically for robotics research in public places similar to the one which has been just published for industry.

REFERENCES

[1] J. Forlizzi and C. DiSalvo, "Service robots in the domestic environment: A study of the roomba vacuum in the home," in *Proc. 1st ACM SIGCHI/SIGART Conf. Human-Robot Interact.*, 2006, pp. 258–265.

[2] F. Faber *et al.*, "The humanoid museum tour guide Robotinho," in *Proc. 18th IEEE Int. Symp. Robot Human Interact. Commun. (RO-MAN)*, Sep./Oct. 2009, pp. 891–896.

[3] P. Ross, "Robot, you can drive my car," *IEEE Spectr.*, vol. 51, no. 6, pp. 60–90, Jun. 2014.

[4] F. Dayoub, T. Morris, B. Upcroft, and P. Corke, "Vision-only autonomous navigation using topometric maps," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst. (IROS)*, Nov. 2013, pp. 1923–1929.

[5] T. Morris, F. Dayoub, P. Corke, G. Wyeth, and B. Upcroft, "Multiple map hypotheses for planning and navigating in non-stationary environments," in *Proc. IEEE Int. Conf. Robot. Autom. (ICRA)*, May/Jun. 2014, pp. 2765–2770.

[6] F. Dayoub, T. Morris, B. Upcroft, and P. Corke, "One robot, eight hours, and twenty four thousand people," in *Proc. Australasian Conf. Robot. Autom.*, Sydney, NSW, Australia, Dec. 2013, pp. 1–7.

[7] (Mar. 2014). *Australian Privacy Act*. [Online]. Available: http://www.oaic.gov.au/privacy/privacy-act/the-privacy-act

[8] M. Quigley *et al.*, "ROS: An open-source robot operating system," in *Proc. ICRA Workshop Open Sour. Softw.*, vol. 3. 2009, pp. 1–5.

[9] G. Grisetti, C. Stachniss, and W. Burgard, "Improved techniques for grid mapping with Rao-Blackwellized particle filters," *IEEE Trans. Robot.*, vol. 23, no. 1, pp. 34–46, Feb. 2007.

[10] D. Fox, "KLD-sampling: Adaptive particle filters," in *Advances in Neural Information Processing Systems*, vol. 14. Cambridge, MA, USA: MIT Press, 2001, pp. 713–720.

[11] E. Marder-Eppstein, E. Berger, T. Foote, B. Gerkey, and K. Konolige, "The office marathon: Robust navigation in an indoor office environment," in *Proc. IEEE Int. Conf. Robot. Autom. (ICRA)*, May 2010, pp. 300–307.

[12] W. Burgard *et al.*, "The interactive museum tour-guide robot," in *Proc. Nat. Conf. Artif. Intell.*, Jul. 1998, pp. 11–18.

[13] S. Thrun *et al.*, "MINERVA: A second-generation museum tour-guide robot," in *Proc. IEEE Int. Conf. Robot. Autom.*, vol. 3. May 1999, pp. 1999–2005.

[14] I. Nourbakhsh, C. Kunz, and T. Willeke, "The mobot museum robot installations: A five year experiment," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst. (IROS)*, vol. 4. Oct. 2003, pp. 3636–3641.

[15] R. Siegwart *et al.*, "Robox at Expo.02: A large-scale installation of personal robots," *Robot. Auto. Syst.*, vol. 42, nos. 3–4, pp. 203–222, Mar. 2003.

[16] H.-M. Gross *et al.*, "TOOMAS: Interactive shopping guide robots in everyday use—Final implementation and experiences from long-term field trials," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst. (IROS)*, Oct. 2009, pp. 2005–2012.

[17] J. Biswas and M. M. Veloso, "Localization and navigation of the CoBots over long-term deployments," *Int. J. Robot. Res.*, vol. 32, no. 14, pp. 1679–1694, Dec. 2013.

[18] *Robots and Robotic Devices—Safety Requirements for Personal Care Robots*, ISO Standard 13482:2014, 2014.

FERAS DAYOUB received bachelor's degree in computer science in 2004, and the Ph.D. degree from the University of Lincoln, in 2012. His Ph.D. thesis was entitled An adaptive spherical view representation for mobile robot navigation in nonstationary environments in 2012. He is currently a Post-Doctoral Research Fellow with the ARC Centre of Excellence for Robotic Vision, Queensland University Technology, Australia, where he is involved in computer vision and mobile robotic applications.

TIMOTHY MORRIS received the B.Math.Sc. and B.Tech. (Hons.) degrees in computing science from the Queensland University of Technology, in 2010 and 2011, respectively. During his time as an undergraduate, he was involved in image processing problems, including a vision sensing hexapod and a shape-based classifier for visual traffic monitoring.

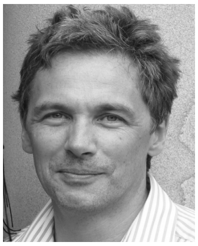

PETER CORKE (F'07) was a Senior Principal Research Scientist with the CSIRO ICT Centre, where he founded and led the Autonomous Systems Laboratory, the Sensors and Sensor Networks Research Theme, and the Sensors and Sensor Networks Transformational Capability Platform. He has held visiting positions with the University of Pennsylvania, the University of Illinois at Urbana—Champaign, the Carnegie-Mellon University Robotics Institute, and Oxford University. He has over 300 publications. He is currently a Professor of Robotics and Control with the Queensland University of Technology, where he leads the ARC Centre of Excellence for Robotic Vision, Australia. He was the Editor-in-Chief of the *IEEE Robotics and Automation Magazine*, a Founding Editor of the *Journal of Field Robotics*, and a member of the Editorial Board of the *International Journal of Robotics Research*, and the *Springer Tracts in Advanced Robotics* series.

• • •