Proceedings of Australasian Conference on Robotics and Automation, 2-4 Dec 2013, University of New South Wales, Sydney Australia

One Robot, Eight Hours, and Twenty Four Thousand People

Feras Dayoub, Timothy Morris, Ben Upcroft and Peter Corke
CyPhy Lab, Queensland University of Technology
firstname.lastname@qut.edu.au *

Abstract

This paper presents an account of an autonomous mobile robot deployment in a densely crowded public event with thousands of people from different age groups attending. The robot operated for eight hours on an open floor surrounded by tables, chairs and massive touchscreen displays. Due to the large number of people who were in close vicinity of the robot, different safety measures were implemented including the use of no-go zones which prevent the robot from blocking emergency exits or moving too close to the display screens. The paper presents the lessons learnt and experiences obtained from this experiment, and provides a discussion about the state of mobile service robots in such crowded environments.

1 Introduction

When moving from datasets and structured static environments to the real world, one of the main challenges facing autonomous robots is the unpredictability of the continuous change in the appearance and structure of their surroundings. This means a set of parameters which produces good results for one scenario inside one environment is not guaranteed to produce the same results when the scenario or the environment is changed.

Current state-of-the-art methods to tackle the problems of mapping, localizing, navigating and planning by mobile robots have been shown to produce desirable and promising results when dealing with each of these problems individually. However, there are few mobile robots which are able to demonstrate long-term autonomy using all these methods together while operating unsupervised in a real-life environment. One example of such an environment is a public event. The main characteristic of these events is that they are densely crowded most of the time which introduces a new set of challenges for mobile robots.

**This research was supported under Australian Research Council's Discovery Projects funding scheme (project number DP110103006).

Figure 1: The Adept Guiabot used during the Robotronica event.

To demonstrate the above points and to highlight the main challenges of operating in a crowded environment, a set of state-of-the-art methods for mapping, localization and path planning has been used to enable a mobile robot to perform a cyclic delivery task with automatic re-charging during a crowded public event. The event was Robotronica, 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.

This paper presents some of the lessons learnt from this experiment which includes the following:

- No personal space: in a crowded environment, people tolerate shared personal space and they expect the same from the robot.
- No-Go zones are an essential requirement for safety purposes: although they appear free in the map, the robot should not execute a plan through certain places (e.g. near emergency exits).

- Strict obstacle avoidance is not always a good thing because it prevents the robot from flowing with the crowd.
- Localization failures are likely: when people surrounding the robot from all directions move with the robot, localization failure is very likely to happen.
- Massive touch screens that operate on infrared (IR) are invisible to IR emitting RGBD cameras.

A video of the robot operating in this environment, and the robot's sensory data, are shown in the video available at `https://www.youtube.com/watch?v=ZJgDB3nu4zs`.

The rest of the paper is organized as follows. In Section 2, we review previous work in deploying autonomous robots in public places. Section 3 provides our experience and the lesson learnt from it. Finally we draw conclusions in Section 4.

2 Robots in the real world

Over the past years, several long-term experiments with mobile robots have been conducted in public places. The earliest examples are the museum tour guides RHINO [Burgard *et al.*, 1998] and MINERVA [Thrun *et al.*, 1999] which appeared in 1998. These robots were equipped with an accurate occupancy grid map of the obstacles in their environments and were required to operate and interact with people visiting the museum. In a confined environment, these robots operated for seven days covering 19 km and two weeks covering 44 km.

Soon after, the Mobot Museum Robot Series [Nourbakhsh *et al.*, 2003] appeared where a number of robots were deployed for five months in a public museum. These robots were able to autonomously navigate for days at a time with automatic re-charging. The Mean Time Between Failure (MTBF) was reported to be between 72 and 216 hours. However, in order to simplify the navigation task and achieve such long-term autonomy, artificial landmarks were installed in the environment and also the existing infrastructure on the ceiling was used for localization with the advantage of being always observable.

Another example of robots operating as tour guides are the RoboX robots [Siegwart *et al.*, 2003]. Around 10 robots were deployed during The Swiss National Exhibition'02 to guide visitors. The MTBF reported from the experiment was 3.26 hours per robot corresponding to approximately 25 human interventions per 10 hours over the whole exhibition. Another example is TOOMAS the Shopping Guide Robot [Gross *et al.*, 2009]. Nine robotic shopping guides were deployed in three different home improvement stores in Germany. TOOMAS uses a state machine to control its operation mode. An occupancy map of the environment was built online while the robot was joysticked through the operation area. The mean run-time was 4.1 hours a day limited by the battery capacity. In 16.3% of cases the robot did not reach its targets, because the hallways were blocked by groups of customers or other dynamic obstacles, or because the robots became lost.

In order to demonstrate a robust autonomous navigation inside an office environment, a long-term experiment using a PR2 robot was conducted by [Marder-Eppstein *et al.*, 2010]. In the experiment, the robot used a prebuilt 2D occupancy of its environment for localization whereas the obstacle avoidance were performed by using a tilting laser range finder. Measurements from the tilting laser finder are used to build a voxel grid encoding information about 3D free, occupied and unknown space. Using the Voxel Grid as an obstacle cost map, the robot demonstrated efficient 3D obstacle detection. The robot operated for 30 hours covering 26.2 miles 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. The authors reported zero human intervention during the course of the experiment. However, when the robot could not reach its current goal, the goal was aborted and the robot moved to another goal.

Another long-term experiment conducted in an office environment used a biologically inspired SLAM system, RatSLAM [Milford and Wyeth, 2010]. RatSLAM is a system which provides a mobile robot with robust autonomous exploration, mapping, localization, and navigation. The system uses an array of sensory inputs which combine vision, laser, sonar, and odometry. In a two week experiment, the system was tested for an office delivery task. Unlike most of the other navigation systems, the robot autonomously explored its environment to build its map instead of being driven by a human. During the two weeks the robot was able to reach its delivery goals with a 99.9% success rate. However, office chairs and glass walls were modified so they were more visible to the sensors.

Long range navigation in outdoor environments has also been attempted by a variety of robotic platforms. One example is Visual teach and repeat robots [Furgale and Barfoot, 2010]. During a learning phase, the rover is driven along a route to capture a set of stereo images. The images are used to build topologically-connected submaps. A path following procedure is done on non-planar terrain, which is handled by switching from localization in three dimensions, to path following in two dimensions using a local ground plane associated with each submap. In one experiment the robot covered a distance of 32 km while autonomous driving in an urban environment.

Another example is the DARPA Challenges for self-driving ground vehicles. Robots such as Junior [Monte-

Figure 2: The Adept Guiabot robot used in this paper. The robot is equipped with laser rangefinder, sonar, stereo vision and RGBD camera. The robot inhabits level 11 of the S-Block in QUT gardens point campus. During operating hours, the robot roams the floor greeting students with its smiley face.

merlo *et al.*, 2008], BOSS [Urmson *et al.*, 2008] and Stanley [Thrun, 2006] have demonstrated robust autonomous navigation through desert and urban areas.

Very recently, the Obelix robot [Kümmerle *et al.*, 2013] demonstrated long range navigation inside a city environment. The robot travelled 3.3 km using local occupancy grids generated from laser scans stored in the nodes of a pose-graph. The pose-graph was created prior to navigation by driving the robot manually. The authors report three instances where human intervention was required and one instance where the emergency stop was activated. The robot used three laser range finders with one mounted downwards to sense the surface directly in front of the robot and the other two were mounted horizontally on the front and back of the robot. GPS was also used for global localization.

In all the above examples the robot was deployed in a relatively open place where it operates around a sparse presence of people. Whereas in this experiment, the robot is deployed in a densely crowded environment which for the best of our knowledge is the first to be tested.

3 Our experience

3.1 The platform

Our experimental platform is a MobileRobots' Research GuiaBot shown in Fig. 2. The sensors used during this experiment were a laser range finder, a RGBD camera (XBOX 360 Kinect) and a sonar. The robot has three on-board computers all running ROS (Robot Operating System).

The GuiaBot normally inhabits a quiet environment on a research floor, level 11 of the S-Block 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. Fig 2 shows the robot in S-Block. People in this environment are robot savvy: academics, postgraduate students and final year students in the robotics, aerospace and Mechatronic disciplines.

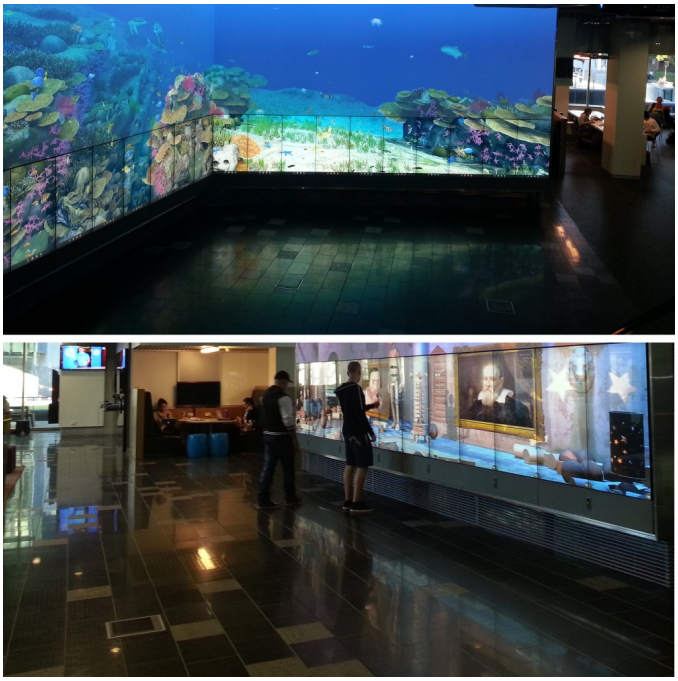

Figure 3: The CUBE, one of the world's largest digital interactive and learning environments with over 40 multi-touch screens. This photo was taken in the early morning before the crowd showed up.

The robot normally serves as a testbed for our work on lifelong autonomous vision-based navigation [Dayoub *et al.*, 2013]. Navigation challenges include low light levels, long sections of texture free walls and glass walls. On average, the robot covers 2.8 km per day.

In order for the robot to be able to operate in this space 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. Our major consideration has been to stop the robot escaping via the lift or via the stairs, the latter being catastrophic for the robot and also very dangerous for anybody on the stairs. We modified the environment with magnetic strips beneath the carpet and retroflective tape on the ceiling in areas around the lift, the stairwell and the firedoors. A hall effect sensor underneath the robot and an industrial

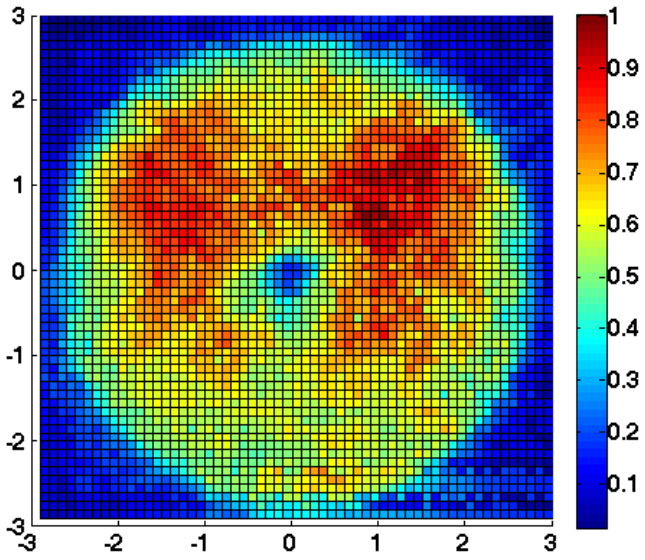

Figure 4: Accumulated obstacles in the 3 m sensory horizon of the robot as a heat map of cells sized $10 \times 10\,\mathrm{cm}^2$. It is centered on the robot's reference frame and orientated such that the horizontal and vertical axes are left to right and behind to front respectively. This shows large activity on the front sides of the robot and lower readings directly in front as obstacle avoidance behavior directs the robot towards less occupied space.

retroreflective tape sensor provide redundant means of emergency stopping the robot, and it requires a key to restart. The robot also has bump sensors at the base and on its sides as well as a number of emergency stop buttons. A large poster in the lift lobby serves to educate people about the robot and the safety issues.

3.2 Robotronica in brief

We had around 8 weeks notice for getting this robot ready for the Robotronica event. We wished to change as little software as possible but the environment had a number of challenges. Fig. 3 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. While the robot deals well containing tables and chairs and glass walled classrooms. 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.

To achieve robust localization and reliable obstacle detection we reactivated the laser scanner and generated a laser based occupancy map of the exhibition floor using GMapping [Grisetti *et al.*, 2007].For localization, the

Figure 5: A view from the camera on-board the robot demonstrating the most common reaction from the public to the RGBD camera and the stereo head. Kinect on the front face of the robot encourages "face-to-face" interaction which inhibits navigation.

robot uses `AMCL`, an adaptive particle filter localizer [Fox, 2001] provided in ROS. The state-of-the-art navigation method used in this paper is provided by the ROS navigation stack (www.ros.org/wiki/navigation). The navigation method 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 [Marder-Eppstein *et al.*, 2010]. Finally, for reliable obstacle detection, the laser scans were augmented by 3D data from a Kinect sensor as described below.

The robot was given a bank of delivery goals and was required to cycle between them. When the battery charge level drops below a predefined threshold, the robot discards its current delivery goal and moves back to the charging station to dock autonomously. During the experiment day the average minimum distance to obstacles was 0.67 m after navigating a distance of 987 m with 293 m resulting from the robot spinning on the spot due to the crowd blocking its path.

The following sections present some of the interesting observations we obtained from deploying the robot in this crowded environment.

3.3 Personal space

It is straight forward to prevent the robot from colliding with a person but it is much harder to stop people from colliding with the robot. This happens especially when the robot is operating in a crowded area. The robot uses sonar to detect people when it backs away from its charging dock, but people are less careful when moving backwards — in one instance on the day a child walked back-

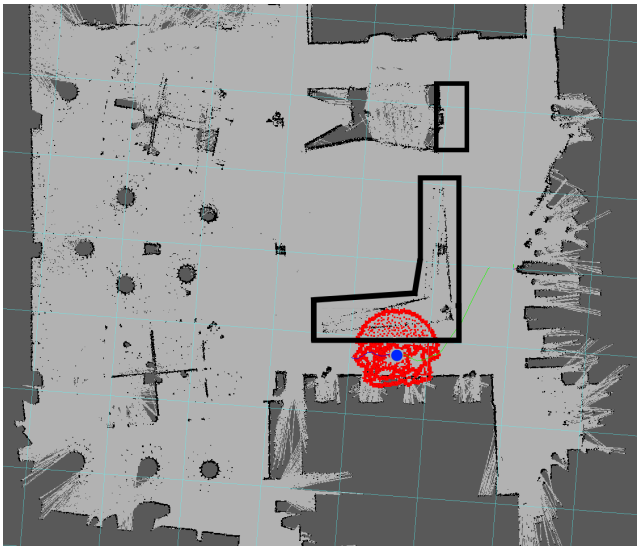

Figure 6: The full map of the floor where the robot has operated. The red region is the obstacles surrounding the robot due to the crowd. The areas contained by the black line are the no-go zones with part appearing as obstacles. The blue dot is the location of the robot. Please see the attached video of the robot in action.

wards into the robot while backing away from a touch screen.

In a crowded environment, people tolerate the invasion of their personal space. According to [Edward, 1966] who introduced the concept of personal space, the average personal space in a western culture is 60 cm on either side, 70 cm in front and 40 cm behind. Fig. 4 shows the accumulated obstacles in a $6 \times 6\,\mathrm{m}^2$ area around the robot expressed as a heat map of cells sized $10 \times 10\,\mathrm{cm}^2$. The figure shows the tendency of people to stand in clusters in front of, and to each side of, the robot and within its personal space.

However people accept that in crowded places this space can be shared and they expect the same would apply to the robot. This means a fixed buffer zone to obstacles in densely crowded environment will slow down the robot and prevent it from flowing with the crowd. Having the robot stop when surrounded is actually unhelpful, in fact in a crowd situation the robot should be attempting to match velocity with the surrounds rather than maintaining free space.

3.4 Human reaction to the robot

We observed fearless interaction of people, especially children, with the robot. Children climbed on the robot and typed on the keyboard. One thing that surprised us was that many people recognized the Kinect sensor on the robot and began gesturing at it, Fig. 5 shows a typical reaction to seeing the Kinect on the robot. All this made navigation at some point very challenging. A close examination of Fig. 4 reveals a slight increase in obstacle concentration in the right front side of the robot compared to the left front side. In fact, in a 0.7 m range around the robot, the right front side was 39.0% occupied on average over the course of the day whereas the left front side was 28.9%. This might be the result of people waving to the Kinect using their right hand which leads them to stand on the right hand side of the robot.

The robot carried a static picture of a smiley face on its forward facing screen. People expected a reaction if they blocked the path of the robot for a relatively prolonged period of time, that is, they expected the face to change expression to upset or angry. We are currently working on making this happen, where the state of the robot will be reflected on the screen by an animated face capable of a limited set of expressions.

3.5 No-Go zones

When deploying the robot in a public place, there are certain areas where the robot should not go for safety and operational reasons. For Robotronica we wanted to ensure the robot kept clear of the touch screen walls to avoid any possibility of damage. A naive solution is to create no-go zones by amending the occupancy map and making these areas 'occupied'. While this prevents the global planner from creating a global path through the no-go zones, the local planner attempts to find shortcuts through unoccupied space and could create a local plan which steers the robot through the no-go zone. Amending the map also has a negative effect on the localization performance since the obstacles in the map are never seen by the sensors.

We found that it was better to cover the no-go zones in the map with virtual obstacles in the form of 3D point clouds and use these clouds as an input to the obstacles cost map. In this case the local planner always steers the robot away from the no-go zones even if the global plan goes through them. Fig. 6 shows part of the no-go zone in the obstacles cost map.

3.6 Effective obstacle avoidance

The environment is problematic due to the presence of many glass walls and the fact that the video walls sit on a cage-like structure which allow laser scans to pass right through. We therefore augmented the laser scan data with an RGBD pseudo scan. This was achieved by was achieved by taking an average horizontal slice out of the RGBD point cloud and selecting the closest points from the laser and the point cloud to create a new scan that was used for obstacle avoidance.

We found that strict obstacle avoidance is not always a good thing because it prevents the robot from flowing with the crowd. One possible solution for such a problem

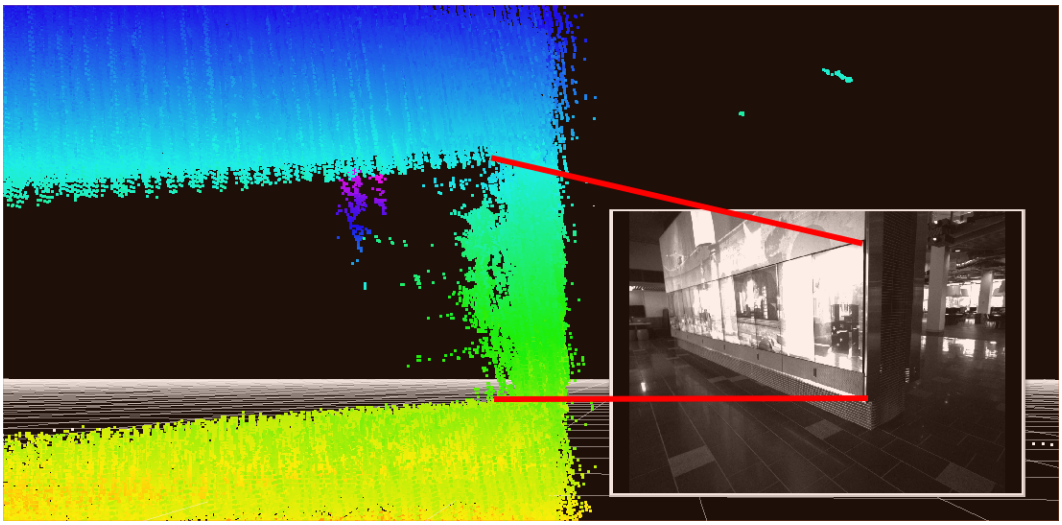

Figure 7: Touch screen that emits infrared are invisible to Kinect.

is to make the range of obstacles a function of the robot's speed and the direction of the movement of obstacles surrounding it. When the robot is moving very slowly with the crowd, a very short range of obstacles is allowed so the robot can move. This means that the robot needs to distinguish between moving obstacles that are in a collision course and obstacles that are flowing with it.

3.7 Localization failures

In a crowded environment there will be prolonged periods of time where the robot is surrounded from directions by people who also move with the robot. This means that localization failure is very likely to happen. One solution to such a problem is the use of landmarks on the ceiling given that the robot is operating indoors and the ceiling has a useful structure for localization.

Although the robot could globally re-localize when out of a crowded region, there is a great risk that it might be in a no-go zone. One possible solution for such a problem is to inflate the virtual obstacles located over the no-go zones according to the uncertainty of the robot pose estimation.

3.8 Touch screens emit infrared light

We found that some touch screens that operate on infrared (IR) are invisible to IR emitting RGBD cameras as shown in Fig. 7. This case re-emphasizes the importance of using passive sensors instead of active ones and demonstrate a case where a stereo vision sensor can outperform an RGBD camera.

4 Conclusions

This paper presented an account of our experiences in deploying an autonomous mobile robot in a densely crowded public event. A laser based occupancy map of the floor of the exhibition was used for localization and path planning. Over the course of an eight hour day the robot undertook a delivery task between a number of locations and recharged itself as required through automatic docking. Over the day the robot was visited by many of the 24,000 people attending the event and operated reliably and safely. Our main observation was that in a densely crowded environment the robot will not have its normal personal space or safety zone. In these situations people use and share their personal space and they expect the same from the robot. Therefore a strict obstacle avoidance strategy can lead to the robot being stationary and safe most of the time. A better strategy is to flow with the crowd and tolerate people in close vicinity when moving at a low speed.

We were also surprised (and heartened) at the expectations of the public to the robot. They are in no way fearful of the robot, they trust it to not collide with them, they expect it to have expressions and emotional states, and to respond to gestures. Hollywood is doing a great job of conditioning people about robots, the challenge for roboticists is to deliver.

References

[Burgard *et al.*, 1998] Wolfram Burgard, Armin B Cremers, Dieter Fox, Dirk Hähnel, Gerhard Lakemeyer, Dirk Schulz, Walter Steiner, and Sebastian Thrun. The interactive museum tour-guide robot. In *Proceedings of the National Conference on Artificial Intelligence*, pages 11–18, 1998.

[Dayoub *et al.*, 2013] Feras Dayoub, Timothy Morris, Ben Upcroft, and Peter Corke. Vision-only autonomous navigation using topometric maps. In *In*

Proc. IEEE/RSJ International Conference on Intelligent Robots and Systems November 3-8, 2013 at Tokyo Big Sight, Japan, 2013.

[Edward, 1966] T Edward. Hall, the hidden dimension, 1966.

[Fox, 2001] Dieter Fox. KLD-sampling: Adaptive particle filters. In *In Advances in Neural Information Processing Systems 14*, pages 713–720. MIT Press, 2001.

[Furgale and Barfoot, 2010] Paul Furgale and Timothy D Barfoot. Visual teach and repeat for long-range rover autonomy. *Journal of Field Robotics*, 27(5):534–560, 2010.

[Grisetti *et al.*, 2007] G. Grisetti, C. Stachniss, and W. Burgard. Improved techniques for grid mapping with rao-blackwellized particle filters. *Robotics, IEEE Transactions on*, 23(1):34 –46, Feb 2007.

[Gross *et al.*, 2009] H-M Gross, H Boehme, Ch Schroeter, St Müller, A Koenig, E Einhorn, Ch Martin, M Merten, and A Bley. Toomas: interactive shopping guide robots in everyday use-final implementation and experiences from long-term field trials. In *Intelligent Robots and Systems, 2009. IROS 2009. IEEE/RSJ International Conference on*, pages 2005–2012. IEEE, 2009.

[Kümmerle *et al.*, 2013] Rainer Kümmerle, Michael Ruhnke, Bastian Steder, Cyrill Stachniss, and Wolfram Burgard. A navigation system for robots operating in crowded urban environments. In *Proc. of the IEEE Int. Conf. on Robotics & Automation (ICRA)*, 2013.

[Marder-Eppstein *et al.*, 2010] Eitan Marder-Eppstein, Eric Berger, Tully Foote, Brian Gerkey, and Kurt Konolige. The office marathon: Robust navigation in an indoor office environment. In *Robotics and Automation (ICRA), 2010 IEEE International Conference on*, pages 300–307. IEEE, 2010.

[Milford and Wyeth, 2010] Michael Milford and Gordon Wyeth. Persistent navigation and mapping using a biologically inspired slam system. *The International Journal of Robotics Research*, 29(9):1131–1153, 2010.

[Montemerlo *et al.*, 2008] Michael Montemerlo, Jan Becker, Suhrid Bhat, Hendrik Dahlkamp, Dmitri Dolgov, Scott Ettinger, Dirk Haehnel, Tim Hilden, Gabe Hoffmann, Burkhard Huhnke, et al. Junior: The stanford entry in the urban challenge. *Journal of Field Robotics*, 25(9):569–597, 2008.

[Nourbakhsh *et al.*, 2003] Illah R Nourbakhsh, Clay Kunz, and Thomas Willeke. The mobot museum robot installations: A five year experiment. In *Intelligent Robots and Systems, 2003.(IROS 2003). Proceedings. 2003 IEEE/RSJ International Conference on*, volume 4, pages 3636–3641. IEEE, 2003.

[Siegwart *et al.*, 2003] Roland Siegwart, Kai O Arras, Samir Bouabdallah, Daniel Burnier, Gilles Froidevaux, Xavier Greppin, Björn Jensen, Antoine Lorotte, Laetitia Mayor, Mathieu Meisser, et al. Robox at expo. 02: A large-scale installation of personal robots. *Robotics and Autonomous Systems*, 42(3):203–222, 2003.

[Thrun *et al.*, 1999] Sebastian Thrun, Maren Bennewitz, Wolfram Burgard, Armin B Cremers, Frank Dellaert, Dieter Fox, Dirk Hahnel, Charles Rosenberg, Nicholas Roy, Jamieson Schulte, et al. Minerva: A second-generation museum tour-guide robot. In *in Proc.International conference on Robotics and automation*, volume 3, 1999.

[Thrun, 2006] Sebastian Thrun. Winning the darpa grand challenge. In *Knowledge Discovery in Databases: PKDD 2006*, pages 4–4. Springer, 2006.

[Urmson *et al.*, 2008] Chris Urmson, Joshua Anhalt, Drew Bagnell, Christopher Baker, Robert Bittner, MN Clark, John Dolan, Dave Duggins, Tugrul Galatali, Chris Geyer, et al. Autonomous driving in urban environments: Boss and the urban challenge. *Journal of Field Robotics*, 25(8):425–466, 2008.