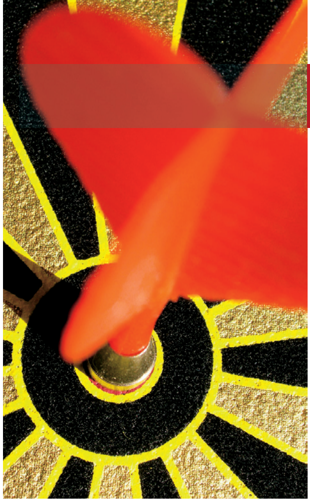

Editor: **Daniel Zeng**, University of Arizona and Chinese Academy of Sciences, zengdaniel@gmail.com

The 1,000-km Challenge: Insights and Quantitative and Qualitative Results

Joydeep Biswas and Manuela Veloso, *Carnegie Mellon University*

In the pursuit of developing autonomous, perpetually deployable service mobile robots, numerous researchers have proposed algorithms for various subproblems, including mapping, localization, navigation, task scheduling, and human interaction. In particular, mapping, localization, and navigation have attracted considerable attention. Although researchers have proposed many such algorithms and map representations, instances of real-world, autonomous long-term deployments of robots to test them in the real world are rare. (See the sidebar, "Related Work in Service Mobile Robots," for examples.) This article presents technical insights and qualitative and quantitative results from extensive real-world deployments of a team of service mobile robots.

We have been working on our service mobile robots, the CoBots,[1] to investigate the challenges in deploying a team of autonomous service mobile robots in a real-world office building. We have used the CoBots to develop and demonstrate several localization algorithms.[2–5] To demonstrate the robustness and accuracy of localization using these contributions, a few years ago, we proposed the 1,000-km Challenge[6]:

> Demonstrate, on a team of deployed autonomous mobile robots, in multiple real-world human environments, the robustness and accuracy in localization over long-term deployments covering a total distance of more than 1,000-km.

The 1,000-km Challenge has the following characteristics:

- robustness and accuracy, to be able to collect quantitative experimental data on the robustness and accuracy of the localization algorithms;
- multiple real-world human environments, to be performed in real environments, so as to expose the localization algorithms to realistic variations;
- long-term deployments, to be performed over deployments spanning multiple years, thus exposing the localization algorithms to environmental variations that would normally occur over such a timespan; and
- team of autonomous robots, to be performed collectively by multiple robots, each with its own unique sensing abilities.

The 1,000-km Challenge began on 17 May 2011 and concluded successfully on 18 November 2014. In this article, we focus on the real-world evaluation of the localization algorithms used on the CoBots over the 1,000-km Challenge and on technical insights into the success of the deployments. Such real-world deployments are invaluable in attesting to the efficacy of the localization algorithms, because they provide experimental results from unstructured, uncontrolled environments. We present quantitative results in two forms: sparse ground truth provided by artificial landmarks placed in the environment and scan matching-based evaluation of localization errors at uniquely identifiable locations. We further present qualitative results of the localization's robustness in terms of the distribution of the logged operator interventions over the 1,000-km Challenge. We augment the quantitative and qualitative results with technical insights into the strengths of the localization algorithms that contributed to the success of the 1,000-km Challenge.

The Scope of Deployments and Data Collected

Since September 2011, four CoBots have been autonomously performing various tasks for users on

1541-1672/16/$33.00 © 2016 IEEE
Published by the IEEE Computer Society

Related Work in Service Mobile Robots

Several research groups have been working on the challenges of continually deployed autonomous service mobile robots. Among them, there are a few instances of deployments in real environments. Shakey was the first robot to actually perform tasks in human environments by decomposing tasks into sequences of actions.[1] Rhino was a robot contender at the 1994 AAAI Robot Competition and Exhibition.[2] Minerva served as a tour guide in a Smithsonian museum.[3] Xavier was deployed in an office building to perform tasks requested by users over the Web.[4] The robots Chips, Sweetlips, and Joe Historybot were deployed as museum tour guides at the Carnegie Museum of Natural History in Pittsburgh.[5] The PR2 robot at Willow Garage has been demonstrated over various milestones in which the robot had to navigate over 42 km and perform several manipulation tasks.[6] The Spatio-Temporal Representation and Activities for Cognitive Control in Long-Term Scenarios (Strands) project held a week-long marathon (http://strands.acin.tuwien.ac.at/marathon.html) to demonstrate continual deployments of a team of robots.[7] YDreams Robotics developed and deployed a team of visitor assistant robots, the Santander Interactive Guest Assistants (www.ydreamsrobotics.com), for the Santander Bank headquarters in Madrid.

In recognition of the challenges of long-term deployments of robots, some researchers have attempted to gather sensor logs from running robots in actual environments over extended periods of time. A five-week experiment at Örebro University collected sensor logs of a robot manually driven around the environment thrice a day, covering 9.6 km and more than 100,000 laser scans.[8] Another dataset, recorded at the University of Lincoln,[9] features omnidirectional images in several settings in an office building. A small subset of the logs analyzed in this article was previously shared in support of the long-term deployments of the CoBots using Corrective Gradient Refinement for localization.[10] Compared to the related work in this area, this article presents results from a fully deployed autonomous robot rather than a manually driven one for the purpose of collecting data.

In contrast to the previous work, our work, to the best of our knowledge, is the first to attempt to provide quantitative measures of accuracy and robustness, along with logged sensor data, while autonomously performing tasks in a real-world environment. We make available the data logs gathered by the robots over the deployments in the hope that they will prove to be useful to other researchers for investigating and testing algorithms for long-term autonomy in real human environments.

References

1. N.J. Nilsson, *Shakey the Robot*, tech. report 323, SRI, 1984.
2. J. Buhmann et al., "The Mobile Robot Rhino," *AI Magazine*, vol. 16, no. 2, 1995, pp. 129–139.
3. S. Thrun et al., "Minerva: A Second-Generation Museum Tour-Guide Robot," *Proc. IEEE Int'l Conf. Robotics and Automation*, vol. 3, 1999; doi:10.1109/ROBOT.1999.770401.
4. S. Koenig and R. Simmons, "Xavier: A Robot Navigation Architecture Based on Partially Observable Markov Decision Process Models," *Artificial Intelligence Based Mobile Robotics: Case Studies of Successful Robot Systems*, MIT Press, 1998, pp. 91–122.
5. I.R. Nourbakhsh, C. Kunz, and T. Willeke, "The Mobot Museum Robot Installations: A Five Year Experiment," *Proc. IEEE Int'l Conf. Intelligent Robots and Systems*, vol. 4, 2003, pp. 3636–3641.
6. A. Oyama et al., "Come On In, Our Community Is Wide Open for Robotics Research!" *Proc. 27th Ann. Conf. Robotics Soc. Japan*, vol. 9, 2009; www.rsj.or.jp/rsj2009.
7. N. Hawes, "Spatio-Temporal Representation and Activities for Cognitive Control in Long-Term Scenarios," 2015; http://strands.acin.tuwien.ac.at.
8. P. Biber and T. Duckett, "Dynamic Maps for Long-Term Operation of Mobile Service Robots," *Proc. Conf. Robotics: Science and Systems*, 2005, pp. 17–24.
9. F. Dayoub, G. Cielniak, and T. Duckett, "Long-Term Experiments with an Adaptive Spherical View Representation for Navigation in Changing Environments," *Robotics and Autonomous Systems*, vol. 59, no. 5, 2011, pp. 285–295.
10. J. Biswas and M. Veloso, "Localization and Navigation of the CoBots over Long-Term Deployments," *Int'l J. Robotics Research*, vol. 32, no. 14, 2013, pp. 1679–1694.

multiple floors of our buildings, including escorting visitors, transporting objects, and engaging in semiautonomous telepresence.[7] The CoBots, shown in Figure 1, vary in their sensing capabilities. CoBots 1, 2, and 3 have a short-range laser rangefinder, the Hokuyo URG-04lx. All the CoBots have one forward-facing depth camera, the Microsoft Kinect. CoBots 2 and 4 also have a second depth camera.

We deployed the robots on several buildings, including the Gates Hillman Center (GHC) and Newell-Simon Hall (NSH) at Carnegie Mellon University (CMU), and the Center for Urban Science and Progress at New York University (NYU). There are 12 floors in total across all the buildings that the CoBots have been deployed on, including floors 3 through 9 in the GHC, 1 through 4 in the NSH, and floor 19 at NYU.

Occupants of the buildings can schedule tasks for the CoBots from our online scheduling interface.[8] Additionally, the robots may also be interrupted by bystanders, who can then schedule tasks on the robot directly using the on-board scheduling interface.[9] The task scheduler assigns and distributes the user-requested tasks among the deployed robots, taking into account transfers between the robots.[10] In addition to user-requested tasks, when the robots do not have any pending tasks, they perform self-assigned tasks commanding the robots to visit randomly chosen locations along the navigable paths in the buildings.

The CoBots navigate through the environment unchaperoned and largely unmonitored. The robots exhibit *symbiotic autonomy*; they autonomously seek human assistance to perform tasks that involve manipulation, because the CoBots do not have arms.[11–13] A robot's execution monitoring scripts track the task execution progress and email the

Figure 1. From left to right, CoBots 1, 2, 3, and 4, which were deployed over the course of the 1,000-km Challenge. The CoBots' sensing capabilities vary, ranging from one or more depth cameras and an optional laser rangefinder.

administrators in the rare cases when it needs assistance. When required, CoBot developers can use the Web-based remote monitoring and telepresence interface of CoBot[14] to inspect the state of the robot and remotely send it commands.

Data Collected

During every deployment, the CoBots log the following data streams:

- drive commands sent to the motors at 20 Hz;
- software exceptions (if any) from all nodes running on the robot;
- operator interventions (if any) using the on-board touchscreen as well as the remote telepresence controls on the website;
- humans detected by the depth cameras at 10 Hz;
- joystick commands received;
- localization estimates at 20 Hz;
- odometry feedback from the wheel encoders at 20 Hz;
- StarGazer observations when received;
- navigation status, including planned path and current command at 20 Hz;
- laser rangefinder scans at 10 Hz;
- obstacle scans computed using the depth cameras at 30 Hz; and
- raw depth images from the depth cameras, at a reduced frame rate of 0:1 Hz.

All data streams are logged at the rates that they are generated, except for the raw depth images, which are logged only at a reduced frame rate of 0:1 Hz to keep the log file sizes manageable.

Table 1 lists the contributions to the 1,000-km Challenge per robot. Figure 2 shows the combined traces of all the locations the CoBots visited over all the maps. The 1,000-km Challenge has resulted in the collection of more than 168 Gbytes of compressed data logs. Table 2 lists the cumulative contents of all the logs. Complete sensor logs collected by the CoBots over the 1,000-km Challenge are available online at www.cs.cmu.edu/~coral/cobot/data.html.

Table 1. Breakup of the 1,000-km Challenge in terms of distances traversed by each CoBot.

Robot	Distance (km)
CoBot 1	36.6
CoBot 2	548.2
CoBot 3	205.1
CoBot 4	216.1

Quantitative Results

We evaluated the error in the localization estimates over the 1,000-km Challenge by two methods: comparison to scan matching and comparison to sparse ground truth.

Accuracy Compared to Scan Matching

Some locations in the environment (for example, corridor intersections) have abundant full-rank long-term map features such that the robot's location can be uniquely determined by scan matching.[15] We call these locations *landmark checkpoints*, and we use them to process the deployment logs offline and estimate the localization errors of the CoBots. The scan matching algorithm is a computationally intensive operation, because it is evaluated over all possible locations and orientations in a 2 m × 2 m search window at a resolution of 0.02 m and 5°. Therefore, the scan matching can only be performed offline for the evaluation of localization accuracy. From the deployment logs, at every instant that a CoBot estimated that it was near a landmark checkpoint, the last observed laser rangefinder scan or depth-image obstacle scan (depending on what was available on that particular CoBot) is used to estimate the instantaneous most probable location of the CoBot by scan matching. When the robots are deployed, they do not slow down or stop specifically at the landmark checkpoints: the online localization and navigation algorithms are even unaware of the landmark checkpoint locations. By comparing the instantaneous most probable location computed by scan matching to the localization estimates of the robot from the deployment log, we estimate the error in localization at that instant.

Table 3 lists the localization errors evaluated by scan matching at landmark checkpoints on the different maps. There are 167 landmark checkpoints in total over all the floors. We

Table 2. Cumulative totals from the data logs collected from all the CoBots over the 1,000-km Challenge.

Property	Value
Duration	1,279.5 hours
Distance traversed	1,006.1 km
Deployments	3,199
Laser rangefinder scans	42,815,389
Depth-camera obstacle scans	54,932,523

omitted maps GHC3, NYU19, and NSH1 through NSH4 because the CoBots encountered insufficient landmark checkpoints on those maps. Figure 3 shows the histograms of errors in localization computed by scan matching at the landmark checkpoints. Figure 4 shows, for each of the different maps, scatter plots of the errors of the robot localization estimates relative to the corresponding landmark checkpoints.

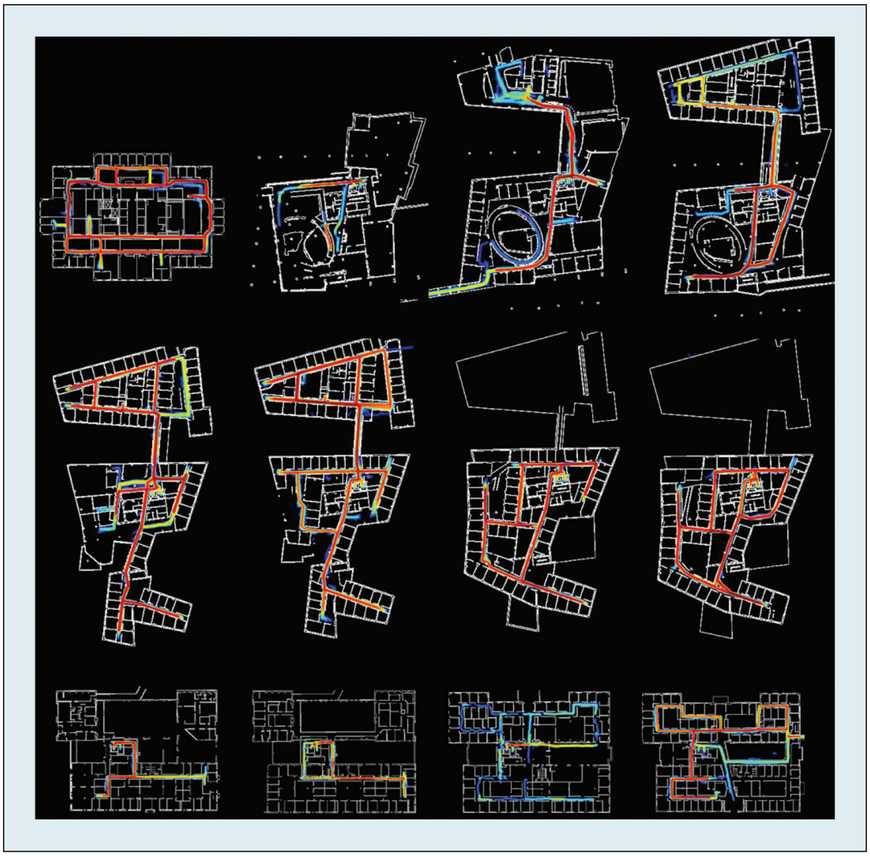

Figure 2. Combined traces of the paths the CoBots traversed over the 1,000-km Challenge: (from left to right, top to bottom) NYU19, GHC3, GHC4, GHC5, GHC6, GHC7, GHC8, GHC9, NSH1, NSH2, NSH3, and NSH4. Locations on the map are color-coded by the frequency of visits, varying from dark blue (least frequently visited) to red (most frequently visited). Because the CoBots visited each floor a different number of times, the color scales are different for each floor.

Accuracy Compared to Sparse Ground Truth

CoBots 2 and 3 are equipped with Hagisonic StarGazer[16] sensors, which detect StarGazer marker patterns mounted on the ceiling and provide relative locations of detected markers with respect to the sensor. Each StarGazer marker consists of four or more retroreflective dots arranged on a 4 × 4 grid. The absence or presence of specific dots on the grid encodes a unique ID number for each marker. We have 46 StarGazer markers placed on the ceiling throughout the GHC to provide sparse ground truth location estimates. We tabulated the markers' locations manually by measuring their locations relative to nearby map features. Although they provide absolute ground truth information, the StarGazer markers have some limitations:

Table 3. Localization accuracy evaluated by scan matching at landmark checkpoints for the different maps over the course of the 1,000-km Challenge.

Map	Samples	Mean error (m)	Median error (m)	Standard deviation (m)
GHC4	294	0.165	0.104	0.159
GHC5	218	0.132	0.101	0.113
GHC6	669	0.105	0.075	0.113
GHC7	6,626	0.132	0.083	0.140
GHC8	527	0.107	0.055	0.145
GHC9	806	0.136	0.080	0.164

- Because the marker locations were tabulated manually, they are subject to human placement and measurement errors.
- The StarGazer sensor reports observed markers at 5 Hz, along with some processing latency. The location estimates of the robot would thus differ when the robot is moving.
- The StarGazer sensor assumes it is always parallel to the markers, but the sensor would in practice tilt with the robot's acceleration and deceleration.

We processed the logs from the CoBots to detect time steps when they detected StarGazer markers, and we estimated the localization errors at those time steps by comparing the location estimates of the robot to the ground truth global locations of the

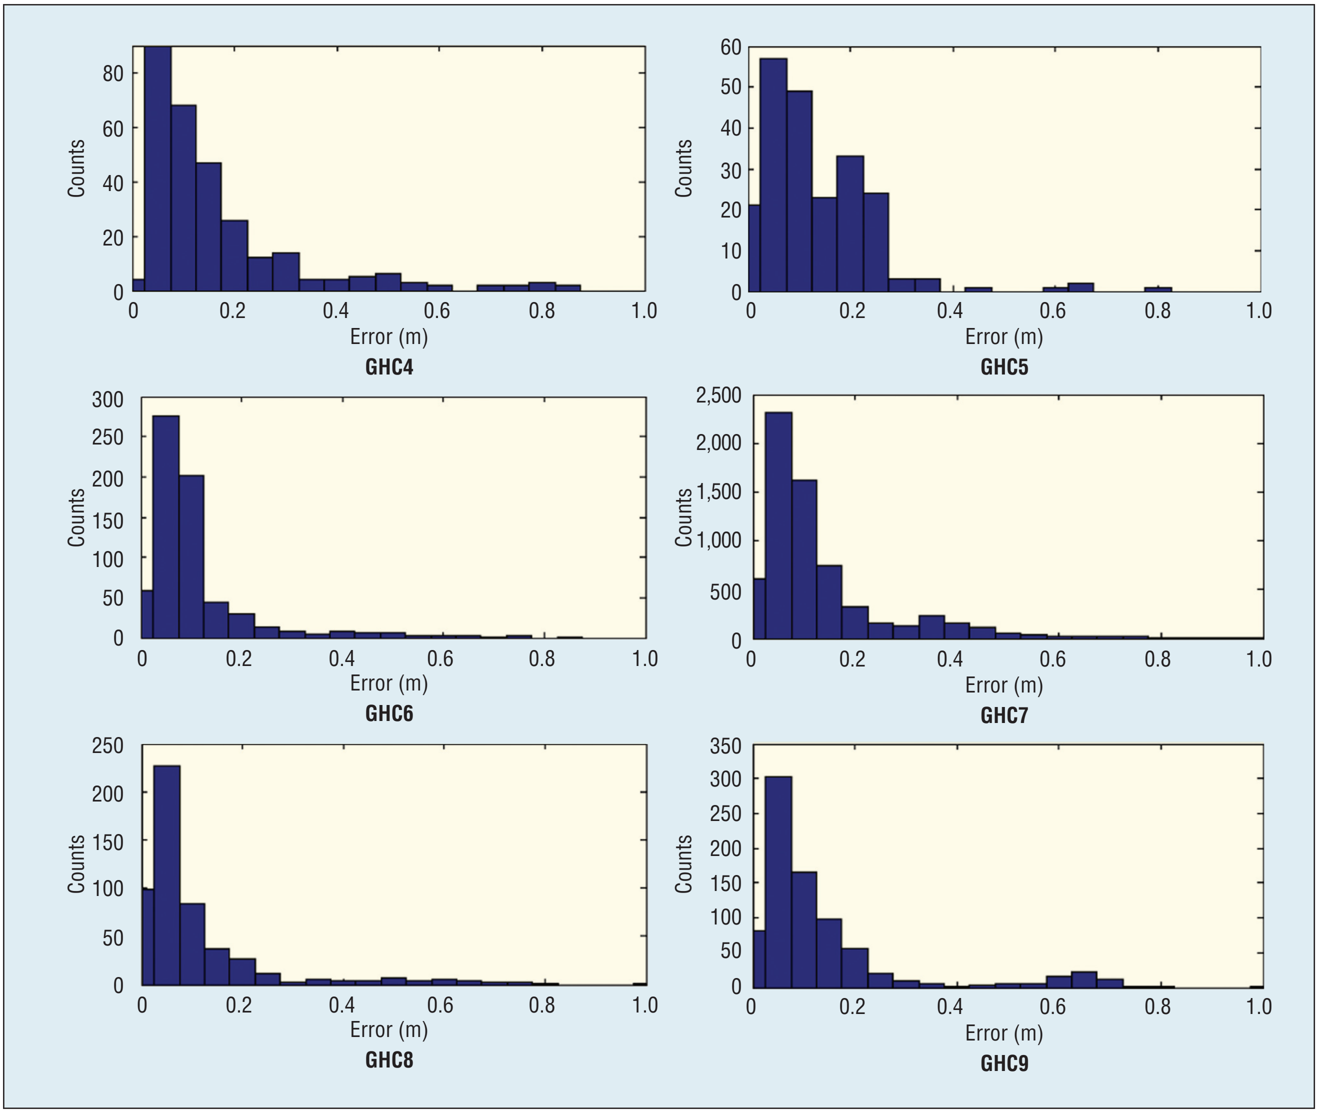

Figure 3. Histogram of scan matching errors on each floor during the 1,000-km Challenge. The histogram distributions indicate that the localization error was less than 0.4 m on all the floors.

observed StarGazer markers. Figures 5 and 6 show the histograms and the scatter plots, respectively, of the localization errors compared to sparse ground truth, and Table 4 lists the mean, median, and standard deviations of errors for the different maps.

Robustness

While deployed, the CoBots log the operator interventions provided. To gauge our localization algorithms' real-world robustness, we count the number of times operator interventions were required to reset localization while the robot was operating autonomously. Table 5 lists the number of times operator interventions were required for all deployments for each map during the 1,000-km Challenge. More than 93 percent of the deployments proceeded without any localization interventions, and there were no deployments with more than three interventions.

Over the duration of the 1,000-km Challenge, we used Corrective Gradient Refinement (CGR) for localization between May 2011 and January 2014,[3] and Episodic Non-Markov Localization (EnML) from February 2014 onward.[5] Table 6 compares the mean distance the robot traversed between interventions, for each map, when using CGR, and when using EnML for localization. The mean distance traversed between interventions is significantly higher for EnML than for CGR, thus demonstrating the former's higher reliability for localization in real-world human environments. We omitted maps GHC3, GHC5, NSH1, NSH2, NSH3, and NYU19 from this analysis because we have insufficient data for both algorithms to compare the mean distance between interventions.

We also logged how far the robots traversed autonomously before operator interventions were required. As

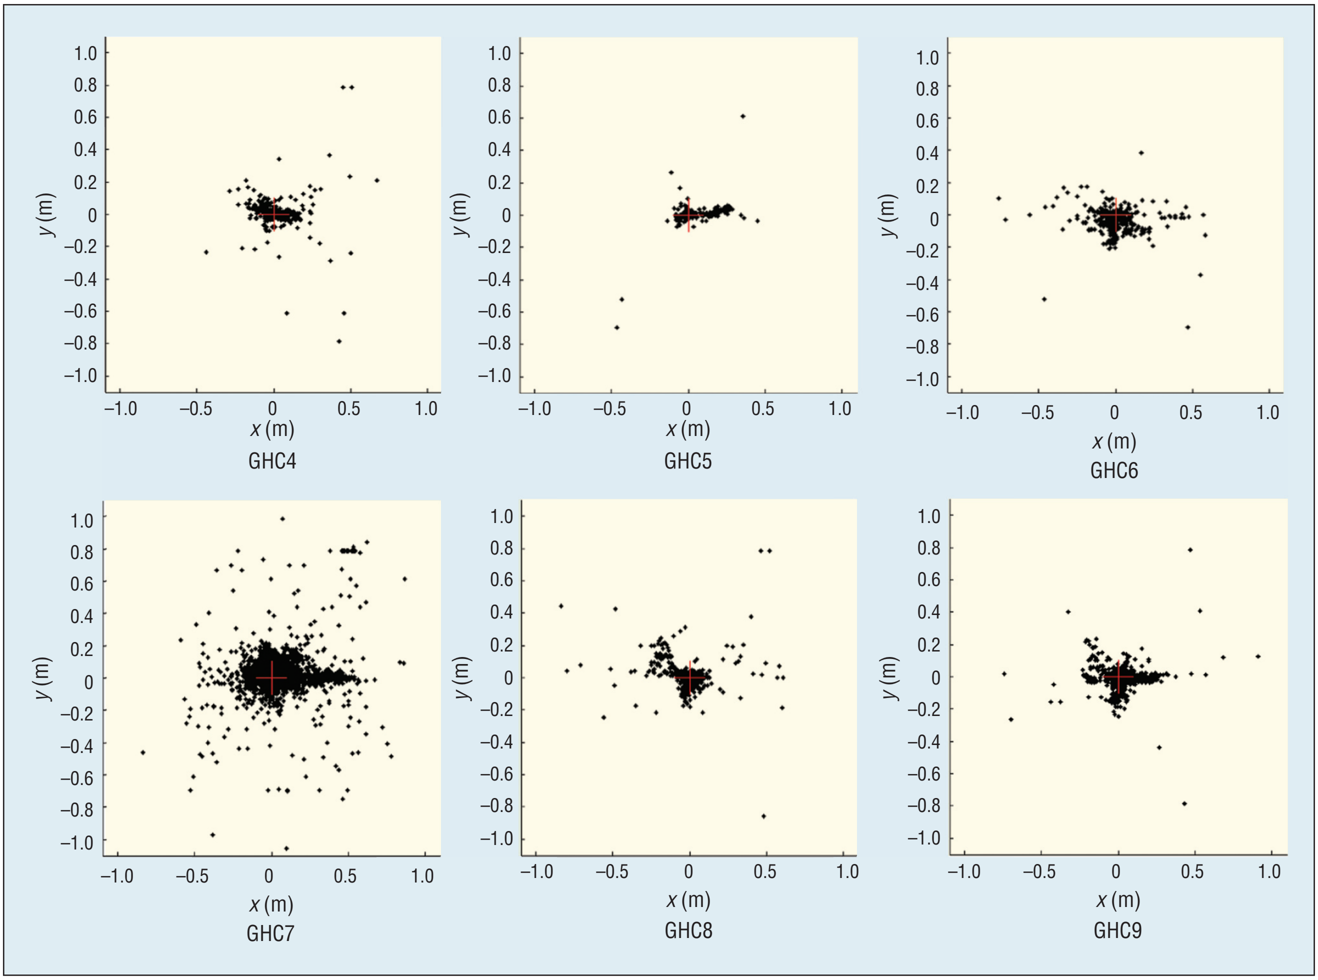

Figure 4. Scatter plots of localization errors evaluated by scan matching during the 1,000-km Challenge. The location bias of the location errors is minimal, but there is some directional bias because of the axis-aligned nature of the corridors on the map.

Table 5 shows, 196 deployments had one or more deployments. We noticed that more than 80 percent of the interventions were required shortly after the robot started moving, within the first 0:5 km. This is largely because the operator provided inaccurate localization initialization at the beginning of the deployments. Furthermore, of the 26 instances where two successive operator interventions were required, more than 75 percent of the successive interventions happened before the robot traversed more than 0:5 km after the first intervention. We believe that this is because when an intervention is required, in many cases it is hard even for the operator to correctly reinitialize the robot's location, thus requiring further interventions shortly thereafter.

Technical Insights

In this section, we provide some technical insights into the features of the localization algorithms that we believe contributed to the CoBots' robustness and accuracy in localization during the 1,000-km Challenge.

Vector Map Representation

The CoBots use a vector map representation of the building's permanent architectural features as a set of line segments.[3] The vector maps are extracted from the building's blueprints and intentionally omit details like the exact locations of movable objects, such as tables and chairs, which are likely to change over time.

Vector maps, unlike the more commonly used occupancy grid maps,[17] are far more compact representations. Furthermore, the accuracy of localization when using an occupancy grid map is limited by the map's cell discretization, whereas localization using vector maps does not suffer from such a limitation. For example, the GHC7 map, which measures 60 m × 112 m, requires 67 kibibytes (KiB) of memory in the vector map format, but would require 255 mebibyte (MiB) of memory with a resolution of 1 cm with 32-bit occupancy values per cell.

Because the vector maps do not include movable objects, the CoBots' localization is unaffected by changes to the poses of the movable objects, such as the closing or opening of doors in the environment.

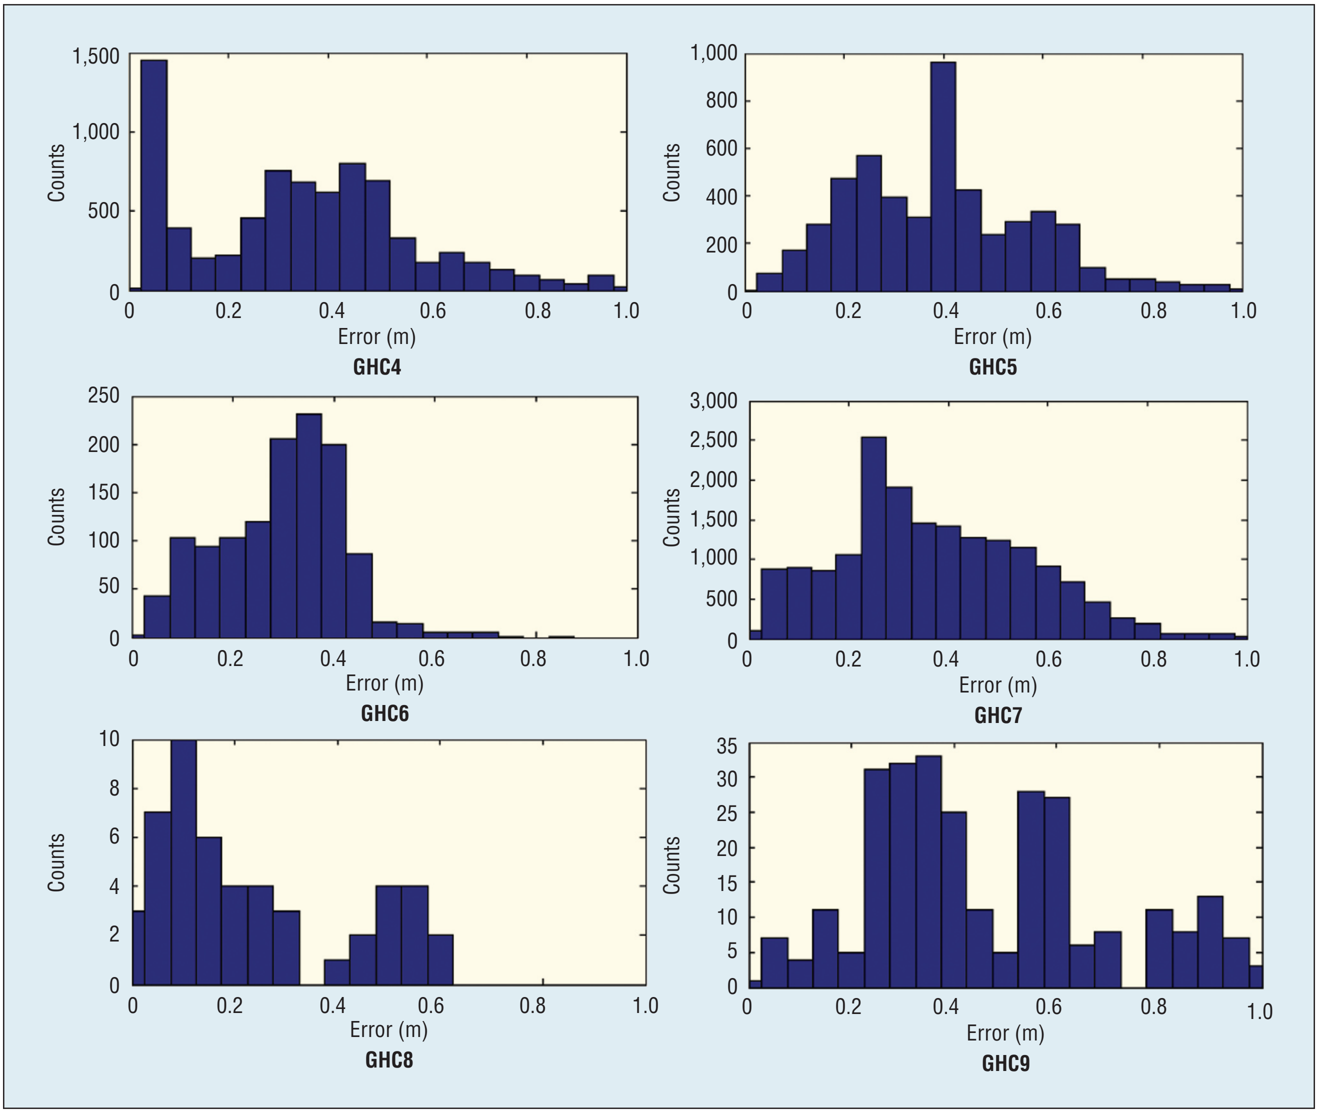

Figure 5. Histogram of errors evaluated by sparse ground truth on each map.

Correspondence matching on the vector maps is performed by analytic ray casting.[4] Once the analytic ray casting is complete, the data association for each observed point (either from a laser scanner or a depth sensor) takes O(1) time. On an occupancy grid map, on the other hand, ray casting for each observed point takes O(L) time, where L is the occupancy grid map's linear size in terms of number of cells. Thus, vector maps allow faster correspondence matching.

Vector maps result in localization that is faster, more memory efficient, and independent of any discretization parameter. Both CGR and EnML benefit from these characteristics because they use vector maps.

Corrective Gradient Refinement

As we empirically demonstrated,[3] CGR has higher accuracy, as well as lower variance across trials than Monte Carlo Localization using Sampling-Importance Resampling (MCL-SIR) when using the same number of particles. This lets the CoBots localize with far fewer particles with CGR than would have been possible with MCL-SIR, while simultaneously requiring lower computational power and providing superior accuracy and robustness.

CGR has the additional benefit of naturally determining, and sampling more from, the degrees of freedom of higher uncertainty. For example, when the robot is traversing down a long hallway, the observations of the two walls would provide ample feedback to uniquely determine the robot's orientation and its location perpendicular to the walls. The robot's location parallel to the walls would be the only uncertain degree of freedom. By refining the particle locations using the state space gradients

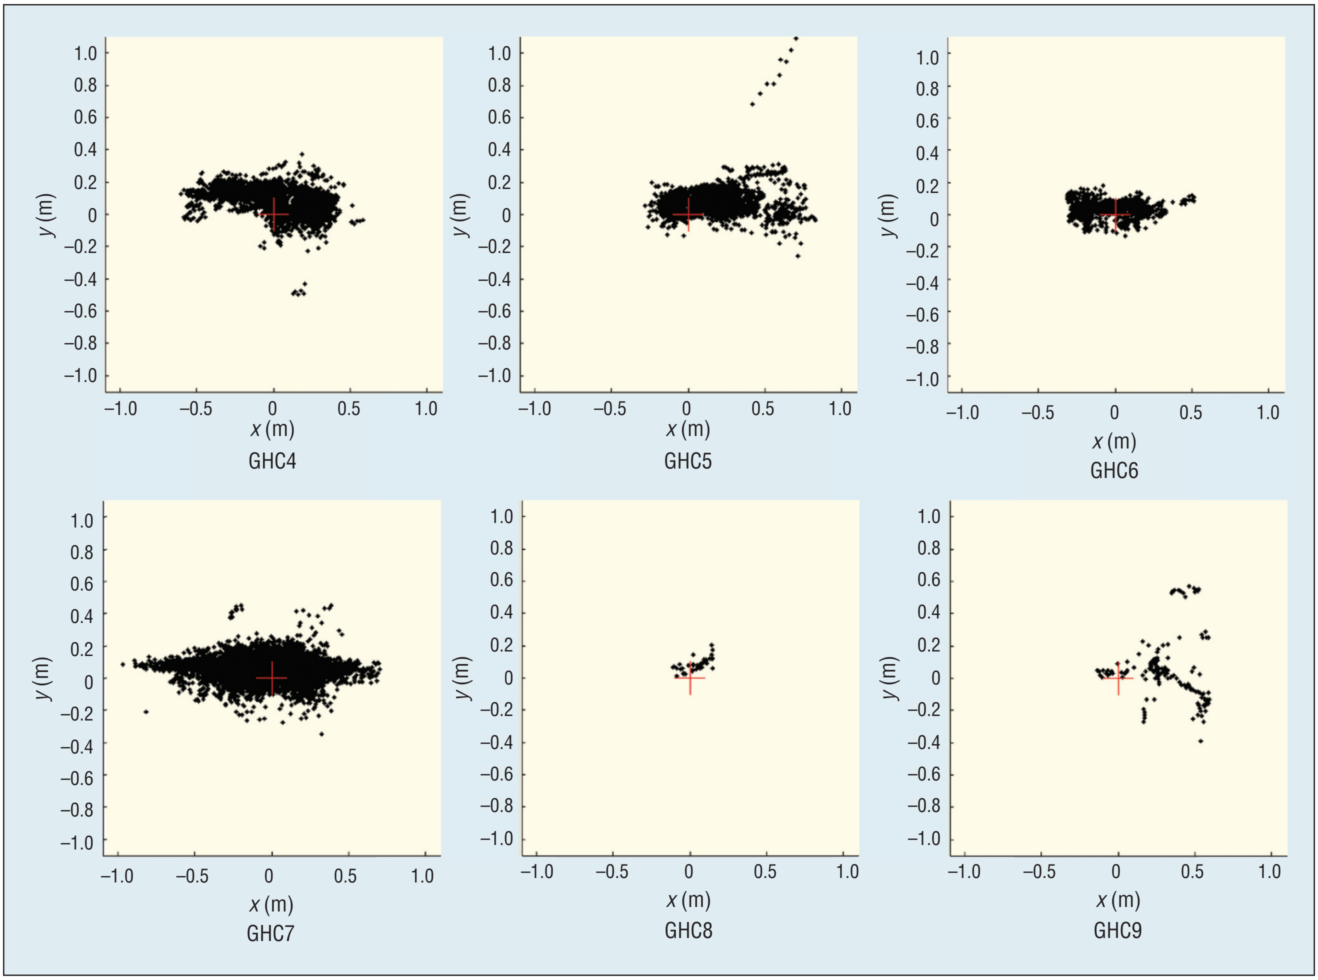

Figure 6. Scatter plots of localization errors evaluated by sparse ground truth during the 1,000-km Challenge.

of the observation likelihood function, CGR correctly distributes the particles more along the direction of the corridor than perpendicular to it.

CoBot 4, being equipped with only a depth sensor, used CGR with the plane-filtered points generated by Fast Sampling Plane Filtering (FSPF)[4] of the observed depth images. Despite the narrower field of view of the depth sensor compared to the laser rangefinder, CoBot 4 still robustly localized using FSPF-CGR, even in the presence of crowds of humans obstructing the robot's sensors. Such robustness is largely attributed to the effective filtering out of nonplanar features such as the humans in the scene, thus minimizing the chances of localization failure due to confusing observations of unmapped objects.

Table 4. Localization accuracy by sparse ground truth for the different maps over the course of the 1,000-km Challenge.

Map	Samples	Mean error (m)	Median error (m)	Standard deviation (m)
GHC4	7,656	0.352	0.343	0.218
GHC5	5,123	0.394	0.400	0.181
GHC6	1,235	0.296	0.314	0.122
GHC7	17,489	0.368	0.341	0.190
GHC8	50	0.234	0.167	0.185
GHC9	276	0.464	0.402	0.232

Episodic Non-Markov Localization

EnML has been used to localize the CoBots since February 2014.[5] Unlike variants of Markov Localization,[18] which ignore observations that do not match the map, EnML explicitly reasons about correlations between robot poses at different time steps because of the observations of unmapped objects. This lets EnML use observations of unmapped static objects such as chairs and tables to provide relative localization feedback while simultaneously

Table 5. Robustness in localization on each map over the 1,000-km Challenge.

Map	Operator interventions per deployment*				
	0	1	2	3	>3
GHC3	14	1	0	0	0
GHC4	126	31	0	1	0
GHC5	44	6	0	0	0
GHC6	290	21	3	1	0
GHC7	2,047	77	9	2	0
GHC8	140	9	1	0	0
GHC9	156	8	0	0	0
NSH1	7	2	0	0	0
NSH2	4	0	1	1	0
NSH3	8	1	0	0	0
NSH4	66	13	0	1	0
NYU19	101	7	0	0	0
Total	3,003	176	14	6	0

*Operator interventions per deployment, listed as the number of deployments with 0, 1, 2, 3, and >3 interventions.

Table 6. Robustness in localization on each map over the 1,000-km Challenge.*

Map	CGR	EnML
GHC4	0.62	4.42
GHC5	—	9.49
GHC6	8.61	9.48
GHC7	5.58	9.02
GHC8	6.04	19.36
GHC9	5.33	20.05
All	4.79	8.13

*Mean distance (in km) traversed between interventions using Corrective Gradient Refinement (CGR) versus Episodic Non-Markov Localization (EnML).

exploiting observations of map features for global localization feedback.

By accounting for and using observations of unmapped movable objects, EnML can robustly and accurately localize the CoBots on challenging floors such as GHC4 and GHC5, in which most of the robots' observations are of unmapped objects, and the unmapped objects occlude the robots' view of map features. Such areas previously were not navigable autonomously; the robots would frequently get lost in them because of insufficient visible map features. At the same time, EnML has dramatically increased the mean distance traversed between operator interventions, from 4:79 km when using CGR to 8:13 km when using EnML (see Table 6).

Sensors

As we mentioned earlier, the CoBots are equipped with different sensor combinations. Owing to the variations in the sensing abilities, the CoBots vary in their ability to handle different types of environments.

With the upward-facing depth camera, CoBot 2 is well-suited for deployments in areas with tall chairs and tables. CoBot 4, with its downward-facing depth camera, can detect small obstacles on the ground, such as laptop power adapters or cables.

The laser rangefinder, with its larger field of view than the depth cameras, makes CoBots 1, 2, and 3 better adapted to deployments in open areas, in which the only observable features are often far off to the side of the robot instead of within the depth cameras' narrow field of view.

Neither the laser rangefinder nor the depth camera sensors on the CoBots are rated for use in direct sunlight. Thus, when the CoBots encounter direct bright sunlight, they occasionally detect false obstacles, and the obstacle-avoidance algorithm brings the robots to a stop.

The sensors used for localization on the CoBots have a limited sensing range and cannot observe the entire length of the hallway that the CoBot is in. Therefore, the localization algorithms must reason about the uncertainty parallel to the direction of the hallway. This is in stark contrast to a scenario in which a robot with a long-range sensor (such as the SICK LMS-200 laser rangefinder, with a maximum range of 80 m) can observe the entire length of every hallway (the longest hallway in the GHC building is about 50 m long), and thus can accurately compute its location with a single reading. The decision to use inexpensive short-range sensors is motivated by cost, because our goal included deploying several CoBots, and to investigate algorithms robust to sensor limitations. Despite the sensor range's limitations, the CoBots repeatedly stop at exactly the right destination locations in front of office doors and always follow the same path down hallways (when there are no obstacles). In fact, in several hallways, the CoBots' repeated traversal along the exact same paths has worn down tracks in the carpets.

Navigation

In our experiences with extended deployment of the CoBots, we learned that a conservative approach to navigation is more reliable in the long term than a more unconstrained and potentially hazardous approach. In particular, the obstacle-avoidance algorithm uses a local greedy planner,[19] which assumes that paths on the navigation graph will always be navigable and can be blocked only by humans. As a result, the planner will not consider an alternative route if a corridor has a lot of human traffic, but will stop before the humans and ask to be excused. Furthermore, because of the virtual corridors, the robot will not seek to side-step obstacles indefinitely. Although this might result in longer stopped times in the presence of significant human traffic, it also ensures that the robot does not run into invisible obstacles in the pursuit of open paths. Although many of the hallways have glass walls, the robot has never come close to hitting them, thanks to its reliable localization and virtual corridors.

System Integration

The CoBots rely on significant automation to ensure continued reliable operation. During deployments, the CoBots are accessible remotely via a telepresence interface that lets the research group examine the state of the robot from the lowest (sensor) levels to the highest (task execution) levels. When a CoBot is blocked for task execution because of lack of human responses to interaction, it automatically sends an email to the research group, mentioning its latest location estimate, task status, and reason for being blocked. Figure 7 shows one such email; the robot had waited for more than 5 minutes for human help at the elevators.

The CoBots' ability to proactively email for assistance when required has been invaluable on a few unexpected occasions. In one such instance, a building occupant blocked the corridor to the robot with cardboard boxes, which resulted in the robot stopping and emailing the developers for assistance. Since then, we have limited the robots' maximum speed while passing by that office, in order to reduce the noise.

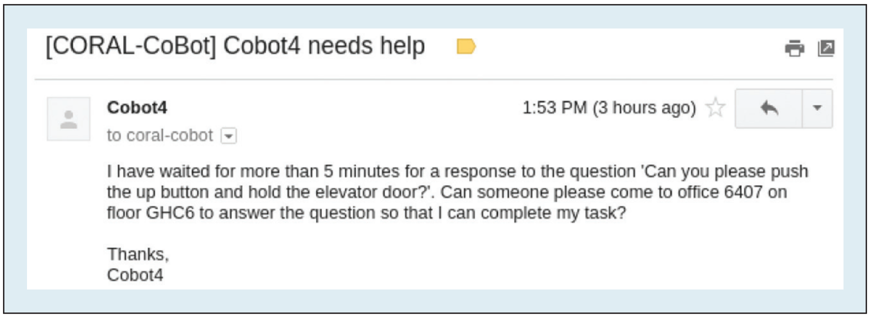
[CORAL-CoBot] Cobot4 needs help

Cobot4 1:53 PM (3 hours ago)
to coral-cobot

I have waited for more than 5 minutes for a response to the question 'Can you please push the up button and hold the elevator door?'. Can someone please come to office 6407 on floor GHC6 to answer the question so that I can complete my task?

Thanks,
Cobot4

Figure 7. An example email from CoBot4, asking for assistance.

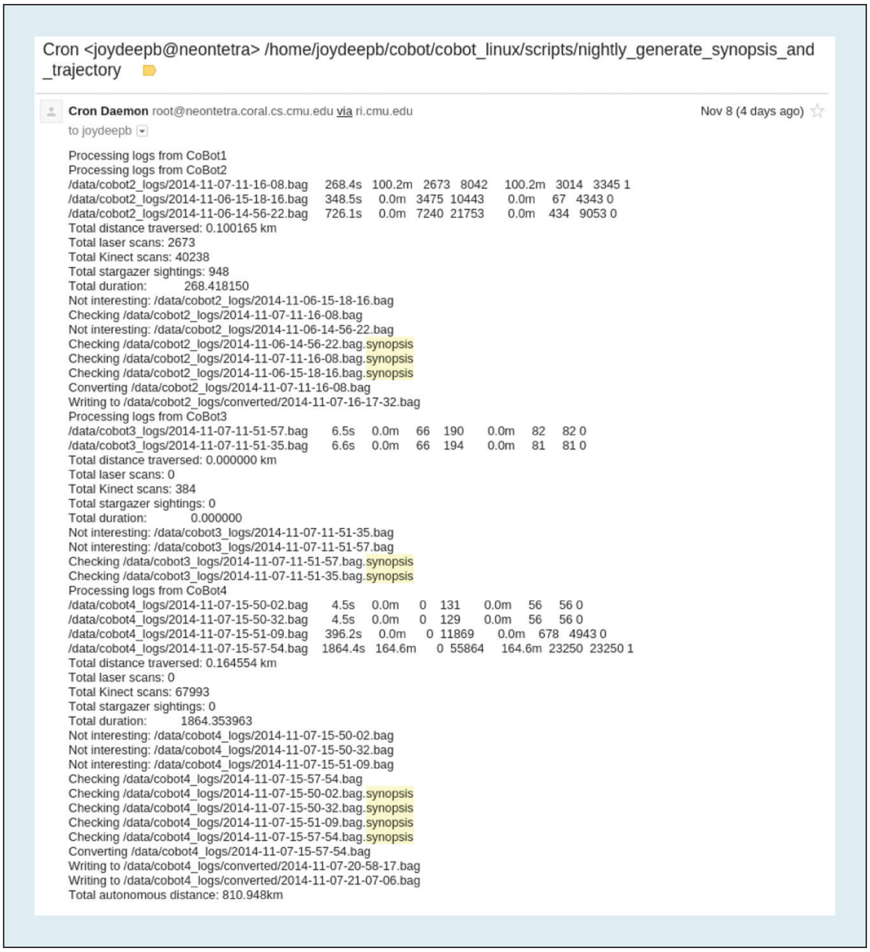
Cron <joydeepb@neontetra> /home/joydeepb/cobot/cobot_linux/scripts/nightly_generate_synopsis_and_trajectory

Cron Daemon root@neontetra.coral.cs.cmu.edu via ri.cmu.edu Nov 8 (4 days ago)
to joydeepb

```
Processing logs from CoBot1
Processing logs from CoBot2
/data/cobot2_logs/2014-11-07-11-16-08.bag   268.4s  100.2m  2673  8042   100.2m  3014  3345 1
/data/cobot2_logs/2014-11-06-15-18-16.bag   348.5s    0.0m  3475 10443     0.0m    67  4343 0
/data/cobot2_logs/2014-11-06-14-56-22.bag   726.1s    0.0m  7240 21753     0.0m   434  9053 0
Total distance traversed: 0.100165 km
Total laser scans: 2673
Total Kinect scans: 40238
Total stargazer sightings: 948
Total duration:        268.418150
Not interesting: /data/cobot2_logs/2014-11-06-15-18-16.bag
Checking /data/cobot2_logs/2014-11-07-11-16-08.bag
Not interesting: /data/cobot2_logs/2014-11-06-14-56-22.bag
Checking /data/cobot2_logs/2014-11-06-14-56-22.bag.synopsis
Checking /data/cobot2_logs/2014-11-07-11-16-08.bag.synopsis
Checking /data/cobot2_logs/2014-11-06-15-18-16.bag.synopsis
Converting /data/cobot2_logs/2014-11-07-11-16-08.bag
Writing to /data/cobot2_logs/converted/2014-11-07-16-17-32.bag
Processing logs from CoBot3
/data/cobot3_logs/2014-11-07-11-51-57.bag     6.5s    0.0m    66   190     0.0m    82    82 0
/data/cobot3_logs/2014-11-07-11-51-35.bag     6.6s    0.0m    66   194     0.0m    81    81 0
Total distance traversed: 0.000000 km
Total laser scans: 0
Total Kinect scans: 384
Total stargazer sightings: 0
Total duration:        0.000000
Not interesting: /data/cobot3_logs/2014-11-07-11-51-35.bag
Not interesting: /data/cobot3_logs/2014-11-07-11-51-57.bag
Checking /data/cobot3_logs/2014-11-07-11-51-57.bag.synopsis
Checking /data/cobot3_logs/2014-11-07-11-51-35.bag.synopsis
Processing logs from CoBot4
/data/cobot4_logs/2014-11-07-15-50-02.bag     4.5s    0.0m     0   131     0.0m    56    56 0
/data/cobot4_logs/2014-11-07-15-50-32.bag     4.5s    0.0m     0   129     0.0m    56    56 0
/data/cobot4_logs/2014-11-07-15-51-09.bag   396.2s    0.0m     0 11869     0.0m   678  4943 0
/data/cobot4_logs/2014-11-07-15-57-54.bag  1864.4s  164.6m     0 55864   164.6m 23250 23250 1
Total distance traversed: 0.164554 km
Total laser scans: 0
Total Kinect scans: 67993
Total stargazer sightings: 0
Total duration:        1864.353963
Not interesting: /data/cobot4_logs/2014-11-07-15-50-02.bag
Not interesting: /data/cobot4_logs/2014-11-07-15-50-32.bag
Not interesting: /data/cobot4_logs/2014-11-07-15-51-09.bag
Checking /data/cobot4_logs/2014-11-07-15-57-54.bag
Checking /data/cobot4_logs/2014-11-07-15-50-02.bag.synopsis
Checking /data/cobot4_logs/2014-11-07-15-50-32.bag.synopsis
Checking /data/cobot4_logs/2014-11-07-15-51-09.bag.synopsis
Checking /data/cobot4_logs/2014-11-07-15-57-54.bag.synopsis
Converting /data/cobot4_logs/2014-11-07-15-57-54.bag
Writing to /data/cobot4_logs/converted/2014-11-07-20-58-17.bag
Writing to /data/cobot4_logs/converted/2014-11-07-21-07-06.bag
Total autonomous distance: 810.948km
```

Figure 8. A synopsis email generated by the nightly processing script running on the central server. The email summarizes the results of processing the deployment logs from the previous day.

Log Management Automation

The CoBots record logs during every deployment, collecting over a gigabyte of sensor data per hour of deployment. The volume of data logs thus generated would understandably be challenging to manage manually. Instead, every robot has a nightly log management script that performs the following tasks at 4:00 a.m. every day:

- compress every data log,
- transfer the compressed data logs to the central server, and
- if the data logs transferred successfully, delete the logs from the robot.

Once the data logs have all been transferred to the central server, they are further processed on the server by a separate log-processing script running on the server. This log-processing script performs the following tasks at 5:00 a.m. every day:

- compute the total autonomous distance traversed for each deployment,
- generate a synopsis image plotting the robot's trajectory over the deployment,
- compute localization errors over the deployments using StarGazer observations and scan matching,
- extract the locations and times of operator interventions, if any, during the deployments,

- convert the data logs into a standard format such that they can be readily shared with researchers,
- generate synopses of the deployments, including the tasks performed and sensor data logged, and
- email a synopsis (see Figure 8) of the deployments to the developer.

Occasionally, a robot might fail to transfer its logs automatically, if it has a poor wireless internet connection. In such a case, an error log is saved locally, and the deployer can check the logs and transfer them manually.

We have presented the results of running the CoBots during the 1,000-km Challenge, spanning deployments over several years in varied environments across multiple floors and buildings. Despite variations and changes in the environments, the robots successfully autonomously traversed all the public areas.

Acknowledgments

This research was supported by NSF award IIS-1012733 and ONR grant number N00014-09-1-1031. The views and conclusions expressed here are the authors' only. We thank the members of the CORAL group, in particular Brian Coltin, Stephanie Rosenthal, and Richard Wang, for the underlying robot task scheduling, human-robot interaction planning, and data-collection deployments. We also thank Mike Licitra for the design and construction of the CoBot robots.

References

1. M. Veloso et al., "CoBots: Robust Symbiotic Autonomous Mobile Service Robots," *Proc. 24th Int'l Conf. Artificial Intelligence,* 2015, pp. 4423–4429.
2. J. Biswas and M. Veloso, "Wifi Localization and Navigation for Autonomous Indoor Mobile Robots," *Proc. IEEE Int'l Conf. Robotics and Automation*, 2010, pp. 4379–4384.
3. J. Biswas, B. Coltin, and M. Veloso, "Corrective Gradient Refinement for Mobile Robot Localization," *Proc. IEEE/RSJ Int'l Conf. Intelligent Robots and Systems*, 2011, pp. 73–78.
4. J. Biswas and M. Veloso, "Depth Camera Based Indoor Mobile Robot Localization and Navigation," *Proc. IEEE Int'l Conf. Robotics and Automation*, 2012, pp. 1697–1702.
5. J. Biswas and M. Veloso, "Episodic Non-Markov Localization: Reasoning about Short-Term and Long-Term Features," *Proc. IEEE Int'l Conf. Robotics and Automation*, 2014, pp. 3969–3974.
6. J. Biswas, *Vector Map-Based, Non-Markov Localization for Long-Term Deployment of Autonomous Mobile Robots,* PhD thesis, Robotics Inst., Carnegie Mellon Univ., 2014.
7. M. Veloso et al., "CoBots: Collaborative Robots Servicing Multi-Floor Buildings," *Proc. IEEE/RSJ Int'l Conf. Intelligent Robots and Systems*, 2012, pp. 5446–5447.
8. B. Coltin, M. Veloso, and R. Ventura, "Dynamic User Task Scheduling for Mobile Robots," *Automated Action Planning for Autonomous Mobile Robots*, 2011.
9. Y. Sun, B. Coltin, and M. Veloso, "Interruptible Autonomy: Towards Dialog-Based Robot Task Management," *Proc. Workshop 27th AAAI Conf. Artificial Intelligence*, 2013; http://brian.coltin.org/pub/sun_aaaiirs_2013.pdf.
10. B. Coltin, *Multi-Agent Pickup and Delivery Planning with Transfers*, PhD thesis, Robotics Inst., Carnegie Mellon Univ., 2014.
11. S. Rosenthal, J. Biswas, and M. Veloso, "An Effective Personal Mobile Robot Agent through Symbiotic Human-Robot Interaction," *Proc. 9th Int'l Conf. Autonomous Agents and Multiagent Systems*, 2010, pp. 915–922.
12. S. Rosenthal and M. Veloso, *Mobile Robot Planning to Seek Help with Spatially-Situated Tasks*, Assoc. Advancement of Artificial Intelligence, vol. 4, 2012, p. 1.
13. S. Rosenthal, *Human-Centered Planning for Effective Task Autonomy,* PhD thesis, Computer Science Dept., Carnegie Mellon Univ., 2012.
14. B. Coltin et al., "Effective Semi-Autonomous Telepresence," *RoboCup 2011: Robot Soccer World Cup XV*, Springer, 2012, pp. 365–376.
15. E.B. Olson, "Real-Time Correlative Scan Matching," *Proc. IEEE Int'l Conf. Robotics and Automation*, 2009, pp. 4387–4393.
16. Hagisonic StarGazer, 2016; http://eng.hagisonic.kr/cnt/prod/prod010102?uid=11&cateID=2.
17. A. Elfes, "Using Occupancy Grids for Mobile Robot Perception and Navigation," *Computer*, vol. 22, no. 6, 1989, pp. 46–57.
18. D. Fox, *Markov Localization: A Probabilistic Framework for Mobile Robot Localization and Navigation*, PhD thesis, Dept. of Computer Science, Univ. of Bonn, Germany, 1998.
19. J. Biswas and M. Veloso, "Localization and Navigation of the CoBots over Long-Term Deployments," *Int'l J. Robotics Research*, vol. 32, no. 14, 2013, pp. 1679–1694.

Joydeep Biswas is an assistant professor in the College of Information and Computer Sciences at the University of Massachusetts, Amherst. He completed the work for this article while at Carnegie Mellon University. Contact him at joydeepb@cs.umass.edu.

Manuela Veloso is the Herbert A. Simon University Professor at Carnegie Mellon University. Contact her at veloso@cs.cmu.edu.