ARMY RESEARCH LABORATORY

Assessment of Navigation Using a Hybrid Cognitive/Metric World Model

by Craig T Lennon, Barry Bodt, Marshal Childers, Bob Dean, Jean Oh, and Chip DiBerardino

ARL-TR-7175 **January 2015**

NOTICES

Disclaimers

The findings in this report are not to be construed as an official Department of the Army position unless so designated by other authorized documents.

Citation of manufacturer's or trade names does not constitute an official endorsement or approval of the use thereof.

Destroy this report when it is no longer needed. Do not return it to the originator.

Army Research Laboratory
Aberdeen Proving Ground, MD 21005-5066

ARL-TR-7175 **January 2015**

Assessment of Navigation Using a Hybrid Cognitive/Metric World Model

Craig T Lennon, Barry Bodt, Marshal Childers, Bob Dean, Jean Oh, and Chip DiBerardino
Vehicle Technology Directorate, ARL

REPORT DOCUMENTATION PAGE	Form Approved OMB No. 0704-0188

1. REPORT DATE *(DD-MM-YYYY)*	2. REPORT TYPE	3. DATES COVERED (From - To)
January 2015	Final	December 2013

4. TITLE AND SUBTITLE
Assessment of Navigation Using a Hybrid Cognitive/Metric World Model

5a. CONTRACT NUMBER

5b. GRANT NUMBER

5c. PROGRAM ELEMENT NUMBER

6. AUTHOR(S)
Craig T Lennon, Barry Bodt, Marshal Childers, Bob Dean, Jean Oh, and Chip DiBerardino

5d. PROJECT NUMBER

5e. TASK NUMBER

5f. WORK UNIT NUMBER

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
US Army Research Laboratory
ATTN: RDRL-VTA
Aberdeen Proving Ground, MD 21005-5066

8. PERFORMING ORGANIZATION REPORT NUMBER
ARL-TR-7175

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)

10. SPONSOR/MONITOR'S ACRONYM(S)

11. SPONSOR/MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION/AVAILABILITY STATEMENT
Approved for public release; distribution is unlimited.

13. SUPPLEMENTARY NOTES

14. ABSTRACT
One goal of the US Army Research Laboratory's Robotic Collaborative Technology Alliance is to develop a cognitive architecture that would allow a robot to operate on both the semantic and metric levels. As such, both symbolic and metric information would be interpreted within the robot's world-model with subsystems providing or interpreting the appropriate type of information. In December 2013, an integrated research assessment was conducted to evaluate progress toward this goal. The robot's ability to receive instructions in structured text, and to interpret those instructions in the context, in both camera and laser detection and ranging images, was applied to the goal of navigating around a building. In this report, the robot's performance is evaluated with 3 measures: 1) grading of the robot's performance during real-time operation, 2) evaluation of planning given the perceptual data within the world model after the testing, and 3) after-the-fact diagnosis of the cause of failed runs by developers of the technology.

15. SUBJECT TERMS
machine intelligence, automated planning, intelligent systems, semantic navigation, world model

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE			Craig T Lennon
Unclassified	Unclassified	Unclassified	UU	74	19b. TELEPHONE NUMBER (*Include area code*) 410-278-9886

Contents

List of Figures

List of Tables

INTENTIONALLY LEFT BLANK.

1. Introduction

The US Army Research Laboratory's (ARL's) Robotics Collaborative Technology Alliance (RCTA) is conducting research to change the relationship between Soldiers and robots. Currently, a robot is a tool that a Soldier uses to accomplish a task. The goal of the RCTA is to develop unmanned ground systems that function as teammates instead of as tools.[1] To do this, the robot must communicate naturally with its human teammates. For example, the robot must accept a command such as "Go around the left side of the building to the back of the building and watch the back door". The robot must be able to do the following in order:

1. Understand the command
2. Identify the relevant building
3. Predict building structure that it cannot see
4. Classify the left side and the back of the building
5. Predict objects referenced in the command when it cannot see them
6. Plan a course around the building
7. Travel its planned course while updating its model of the world with new information available through its sensors

The first step in this process requires issuing a command in language that the robot can understand. The command is conveyed as a highly structured tactical behavior specification (TBS), an ordered set of parameters that have pre-specified values. The parameters and their possible values, as integrated into the robot at the time of the experiment, are described in Table 1. The parameters of the TBS are matched to actions that the robotic platform can execute. An example of a TBS used in the experiment is "Navigate left of the building to a cone near the vehicle". In this example, the command verb is "navigate", the referred object is a building, and the spatial constraint for navigation relative to the building is "left". This conveys to the robot that navigation should be relative to the building and that it should stay on the left side. The target object is a "cone", so it should stop close to a cone. The restriction "near the vehicle" includes a second referred object ("vehicle") and an orientation constraint on the cone relative to that object ("near"). Thus the robot should only stop when it has moved close to a cone that is near a vehicle.

Table 1 Aspects of a tactical behavior specification

Parameter Description	Possible Values
Command verb	Navigate
Command mode	Covertly (optional)
Spatial constraint	Back, side, front
Referred object	Building, vehicle
Target object	Cone, vehicle
Orientation constraint	Left, right, behind, to, near

In all instances throughout this report, the term "cone" actually refers to a large orange traffic barrel used as a navigation target. While the robot is planning based on the TBS, the perceptual system of the robot has detected relevant objects in the environment, which are remembered in the robot's world model.[2,3] Some of these objects have semantic labels, such as vehicles and buildings, and some are simply obstacles around which the robot must navigate. The robot queries the world model for objects matching the parameters of the TBS and grounds the parameter of the TBS to an object within the world model. Some parameters, such as the referred object, which the robot must navigate relative to, must be in the world model before the robot will begin movement, but the exact location of the target object does not need to be known, as the navigation planner can operate on hypothesized objects as well as detected ones. The success of the robot in executing the task described in the TBS depends not only on the efficacy of the planning behavior, but also on the fidelity of the world model, which itself relies on semantic perception (assessed recently in Lennon et al.[4]) and on aspects of the intelligence that hypothesize objects based on partial information.

Some aspects of this intelligence structure were assessed during 16–18 December 2013 at Fort Indiantown Gap, PA, during Integrated Research Assessment 5 (IRA 5). In the next section of this report, the experiments for IRA 5 are presented along with the evaluation criteria used in the assessment. In Section 3, we summarize our evaluation of the robot over 20 of the experimental trials, and in Section 4 we present the conclusions we have reached based on this assessment. Finally, the Appendix contains detailed assessments of the 20 runs analyzed in this report.

2. Experimental Conditions and Assessment Criteria

The intent of IRA 5 was to assess how the intelligence structure performed semantic navigation. As described in Section 1, this performance depends on the TBS as well as on detecting and identifying objects in the environment, hypothesizing undetected objects, and planning paths relative to objects while following navigation constraints given in the TBS. We first describe the experimental conditions under which the robot was tested and then the methods used to assess performance.

2.1 Experimental Conditions

Two buildings at the Fort Indiantown Gap, PA, Combined Arms Collective Training Facility (CACTF) were used as test sites. The first was the church, which is isolated from one visual perspective and can be reliably identified as a building by the semantic labeling system.[4] The church is shown in Fig. 1 and the robot is shown facing the church in Fig. 2. The second venue was the gas station, a cluttered area posing challenges to operation from the perspectives of obstacle avoidance and semantic interpretation. The robot is shown near the gas station in Fig. 3. Figures 2 and 3 also show the large orange traffic cones.

Fig. 1 Back of the church

Fig. 2 Robot in starting position in front of church

Fig. 3 Robot navigating near gas station

The robot used was a Clearpath Husky equipped with an ADONIS infrared camera for semantic object detection and a Micro-LADAR (laser detection and ranging) for obstacle avoidance and semantic object detection. The experiment was planned as a series of vignettes, each consisting of an arrangement of detectable objects placed around either the church or the gas station. Seven vignettes were developed to be used as blocks in the experimental design. The church was the site of vignettes 1 and 4 and the gas station was part of vignette 7. Twenty trials were conducted around the church. In vignette 1, the robot was placed in front of the church with 3 cones placed near and to the right and left of the church. It was then given a TBS directing it to navigate to a cone near, to the left, or to the right of the building, and was expected to navigate to within 1 m of the appropriate cone. The template for vignette 1 is shown in Fig. 4. The template shows the conceptual layout; individual cones—the starting position of the robot—were moved to assess the variability of performance. Multiple cones and varied positions relative to the church reference provided challenges for both semantic labeling and navigation within the TBS instruction.

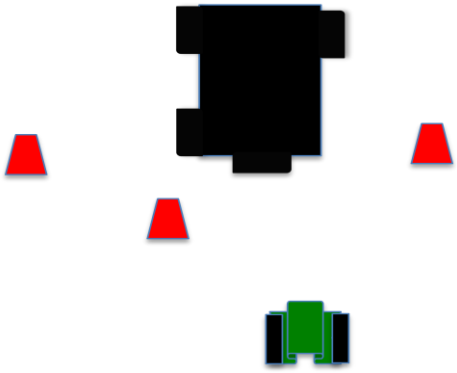

Fig. 4 Template for vignette 1, with the church in black, cones in orange, and robot in green

In vignette 4 (Fig. 5), a vehicle and cones were placed around the church and the robot was given instructions to navigate with respect to these objects and with respect to the church. For example, one TBS was “Navigate left of the building to the cone near the vehicle”, and another was “Navigate left of the building to the cone behind the building”. The robot would then be expected to go to the appropriate side of the building and stop within 1 m of the correct cone. In this simple layout, a few challenges are presented. To begin, the robot is likely to see 2 cones, one left and one right. If the robot is positioned slightly closer to the rightmost cone, as it was for some runs, the TBS “go left” aspect has to add sufficient cost to make the closer, right-most cone less attractive. Similarly, once moving left, the TBS “near vehicle” aspect has to be sufficiently compelling to cause the robot to pass by the first cone on its way to the second. Finally, the TBS sometimes called for a “cone behind the building” as opposed to a “cone near a vehicle”. Goals later in the run remained hypothesized objects until within sensor view.

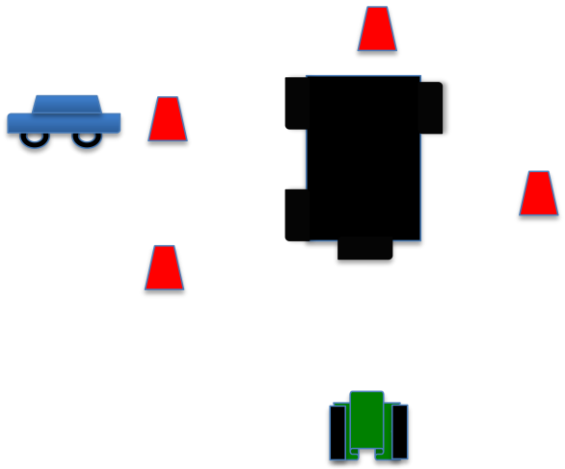

Fig. 5 Template for vignette 4

All the data were used by RCTA researchers for improving their own research, but the runs around the gas station were not sufficient in number or consistency of performance on which to base an analysis. Technical difficulties with the integration prevented us from adhering to the experimental design within each of the vignettes, and some trials needed to be repeated outside of the design of the experiment. When this happened, these runs were designated with a letter. For example, vignette 4, path 1b, is a repetition of the first trial in vignette 4. Some of these technical difficulties had to do with system integration but others were caused by the presence of a substantial amount of snow, which presented difficulties for semantic labeling and caused mobility and navigation challenges. For completeness, the template for vignette 7 is shown in Fig. 6 even though the data for that vignette is not analyzed in this report. Black objects are buildings and curbs while the yellow and gray objects represent concrete pillars around 2 gas pumps.

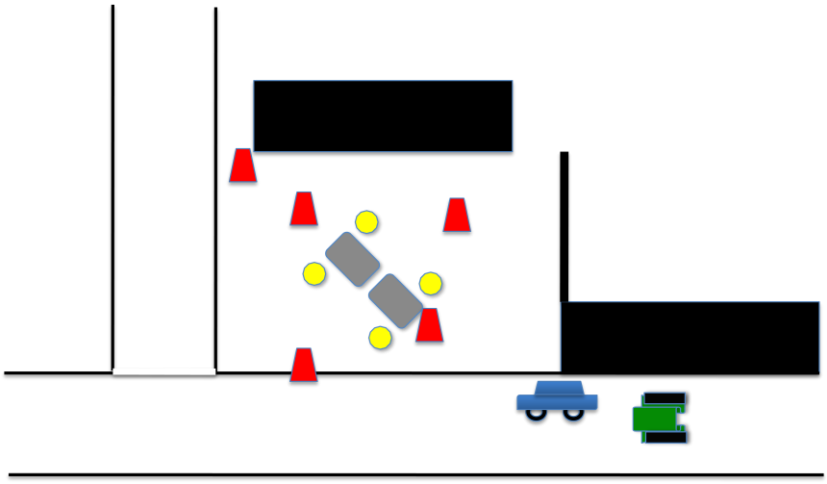

Fig. 6 Template for vignette 7

2.2 Assessment Criteria

Given the variation in vignette conditions and TBS, our intent is to assess whether the system successfully navigated to a position near the target cone while following TBS constraints. We do so by comparing expected behavior to actual behavior as observed by a human. Since a failure to appropriately navigate or get close to the target could be caused my many things, we will also examine aspects of how the intelligence architecture 1) identified objects in the environment, 2) hypothesized undetected objects, 3) recognized its own position, and 4) planned appropriate navigation given TBS constraints.

These 4 aspects will be examined by recreating the robot's model of the world, considering its performance given that model, and including an after-the-fact evaluation of performance by the researchers. We do not attempt to infer whether or not failure to perform as expected was attributable to one of the 4 aspects of the intelligence architecture or to something else. Instead, we provide multiple vantage points of the robot's performance to provide researchers with a richer body of information upon which to base their own conclusions. We begin by describing our method for assessing whether the system successfully navigated to a position near the target cone while following TBS constraints.

For each of the 20 runs, a human observer sketched the robot's path over the course of the run as well as the positions of the objects that were part of the vignette. The observer also evaluated the performance of the robot according to the criteria shown in Table 2.

Table 2 Human observer scoring for vignette 1, paths 2–4

Aspect	Assessment	Comments
Appropriate navigation	Partial success	Command: "Navigate right of building to a cone right of the building." Started right, turned left, then back right
Following constraints	Success	. . .
Reaching goal	Success	Stopped within 1 m of cone C3
Other	. . .	. . .
Overall	P 1[a]	. . .

[a] See Table 3

The robot was evaluated based on how appropriately it navigated to its goal, how well it followed constraints, and how close it got to the goal. It each category, it was rated as success, partial success, or failure. Based on successes and failures, it received a code summarizing its overall performance according to the criteria listed in Table 3.

Table 3 Overall performance evaluation codes

Overall Performance	Code
Success at all aspects	S
Partial in one aspect, success in others	P 1
Partial in 2 or more aspects but no failures	P 2
Failure in one (goal/miscellaneous) but success in others	F 1
Failure in one (goal/miscellaneous) and partials in others	F 2
Failure in 2 components	F 3

As the robot is not permitted to use an a priori map, it must create its model of the world anew with each trial of the experiment. At the end of each trial, this world model was saved, and the robot's experience during the trial was later reconstructed to assess the robot's performance in the 4 aspects of the intelligence architecture listed at the beginning of this subsection. For example, consider how the robot identified objects in its environment. This relies on semantic perception correctly labeling the object and on intelligence to determine when it is seeing the same object it saw before. We do not attempt to disentangle the performance of these 2 subsystems, but we do count the total number of cones and vehicles that were identified during the run and plot the locations at which they were perceived. Researchers can use this data to explore which subsystem contributed to over- or under-identification of semantically labeled objects.

The second aspect we consider is the hypothesizing of objects. As part of the intelligence architecture, ACT-R[5] software was used to hypothesize the undetected walls of buildings based on the observed walls. It was provided with a priori knowledge of the shape and dimensions of the church and other buildings in the CACTF but had to determine, based on sensed information, that the church was the building being viewed and, finally, its orientation and location. The

positions chosen by ACT-R were examined along with information available in the world model, such as the positions of walls and the path the robot had taken. Then an integration and assessment team member assessed whether the building prediction was reasonable given only the information available within the world model.

The assessment of how well the robot recognized its own position was a subjective one, made by comparing the sketch of the robot's path made by a human observer during the trial with the robot's perception of its path as recorded in its map of the world in the world model. This assessment constitutes a qualitative judgment as to whether the robot made an appropriate choice given the information in the world model. The assessment of the plan given the robot's world model need not be the same as the human observer's assessment of navigation as executed, which was made by observing the movement of the robot in the real world.

The assessments of these 4 aspects are included as tables in the Appendix, and an example is included as Table 4. Columns indicate which aspect was assessed, the time of the run at which the observed activity occurs, an aspect dependent code, and comments. The Time column is either a time or is Final, meaning the observation is made after all activity in the run is complete. The values of the Assessed column and the codes for each assessed aspect need more explanation. These assessment codes are listed in Table 5.

Table 4 Example assessment of vignette 1, path 2-4 based on world-model data

Assessed	Time	Code	Comments
Building	14:39:10.5	W, N	Fig. 8
Vehicles	14:39:10.5	. . .	7 observed at this point, Fig. 8
Cones	14:39:10.5	. . .	7 observed at this point, Fig. 8
Navigation plan	14:38:03.8	R	It is hard to determine which is the right side of the building in Fig. 9 left, but the plan is appropriate when the building is correctly predicted in Fig. 9 right.

Table 5 Meaning of assessment codes

Assessed	Code	Meaning
Building	W	Prediction not consistent with walls as registered in world model
Building	N	Prediction inconsistent with navigation information
Nav. plan	C	Violated navigation constraints
All	O	Other: see comments in the Table 4
All	R	Reasonable action or prediction

An assessed column value of Building indicates an assessment of the robot's performance with respect to hypothesizing the undetected parts of a building. The assessment is based on consistency with other world model data and not on how well the hypothesis corresponded to ground truth. A W code indicates that the building prediction was not consistent with the position

of walls as registered in the world model. An N code indicates inconsistency with navigation information in the world model—in particular, that the building was predicted as being on top of the robot or in a position in which the robot had recently traveled. The R code indicates reasonable predictions, i.e., if they generally correspond to the position of the known walls without being inconsistent with navigation information. The robot may have many Building predictions throughout a run, and all must be reasonable for an R code. The Comments column indicates the figure to which one might refer to see why the code was chosen, as well as other relevant comments. For all assessments, the O code stands for other comment. For an example of a Building assessment, consider the snapshot of the world model shown in Fig. 7: The walls are shown in black, the route the robot has traveled is dark green, cones are dark yellow, and vehicles are purple. All of the preceding objects are plotted as registered in the world model, thus this is how the robot thinks the world looks. The robot has just re-hypothesized the position of the building. Its old hypothesis was that the building was in the position shown by the blue dashed outline, and its new hypothesis is that the building is in the position shown by the red solid outline. Neither is consistent with the registered positions of the walls, and the red outline is on top of where the robot has traveled. If only the blue outline was predicted, the code would be W and the code for the red outline would be W and N. The time at which this re-prediction occurred is shown at the top of the figure down to tenths of seconds.

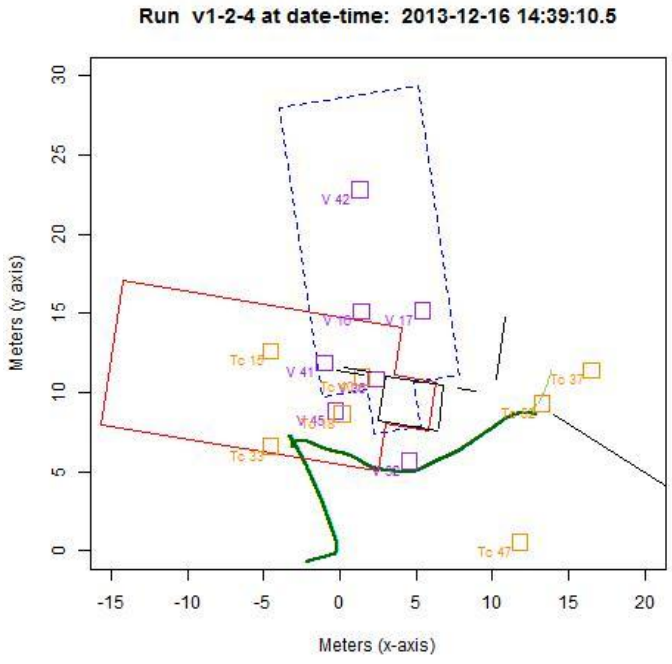

Fig. 7 Example of a building transformation in vignette 1, run 2-4

If the assessed column contains “vehicle”, “cones”, or “fire hydrant”, the aspect assessed is the identification of objects in its environment. Generally there will be no code but rather just a count of how many of each object is registered in the world model. The actual number of cones present is known for each run, as it was part of the vignette. In this run, there were 3 cones actually present. In some vignettes, one car was part of the vignette, and in others there were none. Regularly, however, there were cars on the street near the church, so there could have been more than one visible to the robot. We have chosen to plot them, list the number observed, and leave it to the researcher to decide whether they think the cars were really present. An inaccurate count of cones and vehicles does not always come from inaccurate semantic perception. A failure of the robot to recognize when a cone is the same one it has seen before can also cause false positives, and if the robot does not navigate near a cone, it may not see it. But the counts and plots do give us insight into the robot’s model of the world and how confusing that world can be to a planning system reasoning upon it.

Navigation planning was assessed on whether the plan followed constraints and was appropriate assuming the objects and building prediction in the world model had been correct. Did the robot plan to go to the side of the building it was told to go to or did it violate these navigation constraints and go in a different direction? A C code indicates that the navigation planner violated one of the constraints given in the command, and R indicates that all planned paths followed the constraints given the robot’s perception of the world. Fig. 8 shows 2 examples of planned paths from vignette 1, run 2-4. The command was “Navigate right of building to a cone right of the building”. The left image is difficult to assess because the robot has hypothesized the building position incorrectly and there are no cones on what it considers the right side. But it did go toward the right side of the hypothesized building and then tried to find a cone. This is the most difficult of the plans to assess. The vast majority of trials have plans that are easy to evaluate, such as the right image of Fig. 8, which is a plan from later in the same trial in which the robot is clearly going to the cone on the right side of the building. Navigation in this trial is evaluated as following constraints reasonably well (R) given the robot’s model of the world at the time the plan was generated. In evaluating navigation, we also take into account that the planner might not receive updated building positions immediately and allow up to 5 s for the planner to alter the plan to a reasonable course after the building has moved.

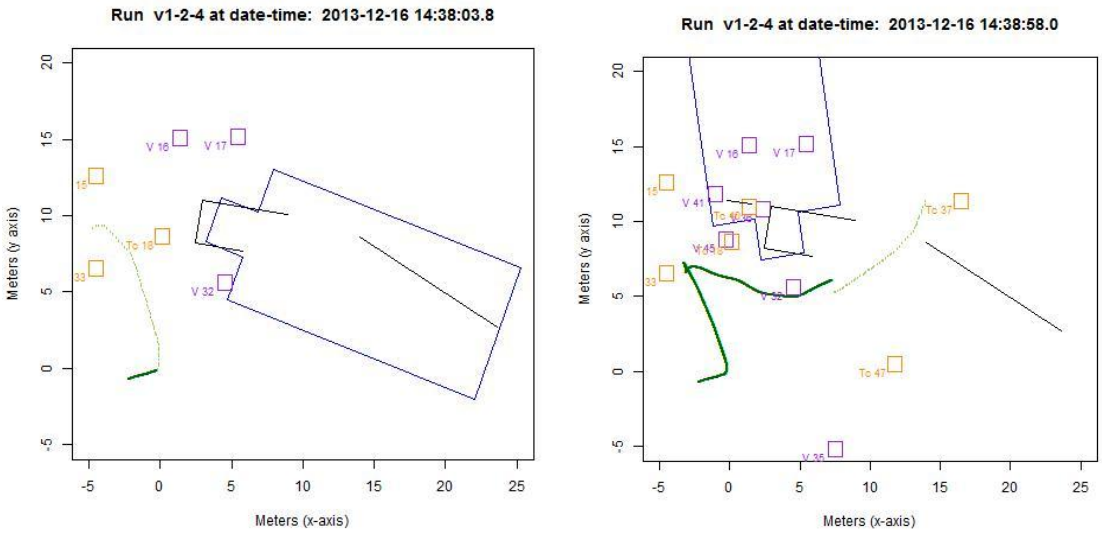

Fig. 8 Example of navigation planning assessment from vignette 1, run 2-4

Localization is assessed as successful as long as the robot's view of where it went roughly corresponded with the human observer's view, graded by comparing his/her sketches with the robot's map after the fact. Localization was generally successful and is only noted in the assessment when it was not successful.

The performance of the robot was also assessed by one of the researchers by assigning error codes, shown in Table 6, to runs based on observation of things that went wrong during the runs. Environmental issues describe problems with driving through snow and over ice. Grounding refers to the robot's interpretation of spatial relationships and mapping specific object instances to the symbols used in a TBS command. Object detection denotes errors like mislabeling of objects, false positives, or false negatives. Platform failure means the robot had a mechanical malfunction. Prediction error indicates an error by the building predictor in predicting the location or orientation of the building.

Table 6 Error codes assigned by the researcher/operator

Code	Description
EN	Environmental issues, e.g., snow
GR	Grounding error
OD	Object detection error
PL	Platform failure
PR	Prediction error

3. Assessment Summary and Analysis

Based on the assessment criteria of Section 2.2, we evaluated the 20 runs around the church and summarized the results in Table 7. The Trial column indicates the vignette and path for the trial. The next 3 columns are the human observer grades for navigation, constraints, and goal, with the fifth column indicating the overall assessment assigned by the human observer. The next 4 columns are the world model-based assessments for the building, navigation, vehicles, and cones categories. The second to last column is the number of cones actually present in the trial, and the final column is the code for the engineering assessment made according to the criteria of Table 5. Localization is not listed here because it was successful in all trials except possibly in vignette 4, path 1b, when the robot got stuck in the snow.

Table 7 Ground truth- and world-model-based assessments of robot performance

Trial	H Nav.	H Cons.	H Goal	H Score	W Bld.	W Nav.	W Veh.	W Cone	Cone	Opr.
V1 P1a	S	S	S	S	R	R	4	1	3	S
V1 P1b	S	S	S	S	W, N	R	6	6	3	S
V1 P24	P	S	S	P 1	W, N	R	12	8	3	GR
V1 P24b	P	F	F	F 3	R	C	4	5	3	GR
V1 P35	P	F	F	F 3	N	O	4	4	3	GR
V1 P35b	P	F	F	F 3	W, N	O	12	6	3	EN, GR
V1 P6	S	S	P	P1	W, N	R	3	3	3	OD
V1 P6b	F	S	F	F 3	W, N	R	3	5	3	PL
V1 P6c	S	S	S	S	R	R	7	4	3	OD
V4 P1a	F	F	F	F 3	W	C	8	2	4	GR
V4 P1b	S	S	P	P 1	R	R	2	2	2	OD
V4 P1c	S	S	S	S	W, N	R, O	0	2	2	OD, GR
V4 P2	P	P	F	F 1	W, N	R	10	2	2	PR
V4 P3	P	F	F	F 3	W	C	3	0	2	GR
V4 P3a	P	P	F	F 1	W, N	R	1	0	2	OD, PR
V4 P3b	P	F	F	F 3	R	C	1	0	2	GR, OD
V4 P3c	S	S	P	P 1	W, N	R	6	3	2	S
V4 P3d	S	S	S	S	W	R	4	4	2	OD
V4 P4a	F	F	F	F 3	W	C	13	1	2	PR, OD
V4 P4b	S	S	S	S	W, N	R	9	4	2	S

Considering Tables 8–12, we can make a few observations. First, localization was generally successful, while building prediction shows room for improvement. This suggests that hypothesizing undetected aspects of buildings might be improved by taking into account navigation information, in particular the path on which the robot has traveled. One could, for example, reject predicted buildings that cover the robot's previous path. By contrast, the number of vehicles shown Table 7, and their locations as plotted in the individual assessments in the Appendix, make vehicle detection seem less reliable, so hypothesizing building positions to

avoid vehicles might not be as successful. More generally, prediction of hypothesized objects should be designed to take into account information in the world model only when the designer regards that information as generally reliable. Localization and wall detection may be more mature technologies and thus might be good for such a purpose.

Table 8 Counts of ratings by for combinations of human assessment criteria

H Nav.	H Cons.	F	P	S
F	F	2	0	0
	P	0	0	0
	S	1	0	0
P	F	5	0	0
	P	2	0	0
	S	0	0	1
S	F	0	0	0
	P	0	0	0
	S	0	3	6

Table 9 Counts of codes from some world model-based evaluations

	W Nav.			
W Bld.	O	C	R	
W, N	1	3	10	14
N	0	0	1	1
R	0	2	3	5
	1	5	14	

Table 10 Counts of codes from world model navigation and operator evaluations

	Operator								
W Nav.	EN, GR	GR	GR, OD	OD	PL	PR	PR, OD	S	
C	0	3	1	0	0	0	1	0	5
O	1	0	0	0	0	0	0	0	1
R	0	2	1	4	1	1	1	4	14
	1	5	2	4	1	1	2	4	

Table 11 Counts of codes from world model building prediction and operator evaluations

	Operator								
W Bld.	EN, GR	GR	GR, OD	OD	PL	PR	PR, OD	S	
N	0	1	0	0	0	0	0	0	1
R	0	1	1	2	0	0	0	1	5
W, N	1	3	1	2	1	1	2	3	14
	1	5	2	4	1	1	2	4	

Table 12 Counts of researcher identified errors occurring over the 20 trials

Code	Description	No.
EN	Environmental issues	1
GR	Grounding error	8
OD	Object detection error	8
PL	Platform failure	1
PR	Prediction error	3

A second observation is that identifying objects and their positions seems to be an issue, in particular the identification of vehicles. Semantic perception might contribute to this but the lack of an ability to identify objects as previously seen is also a suspect. The robot will eventually need a process for deciding when an object is something it has seen before.

Finally, navigation planning was inappropriate or violated constraints in 5 out of the 20 trials considered. Even if the robot eventually gets to its goal, violating navigation constraints will strain the trust of the operator. This might be a good place for human robot interaction to become involved, with a robot showing an intended path and asking if it can take that path before execution.

4. Conclusions and Recommendations

IRA 5 was the last major experiment prior to the RCTA capstone study scheduled for Fall 2014. In evidence were the difficulties of the RCTA challenges of building prediction, semantic perception, and semantic navigation, each proving substantial from both software and hardware perspectives. It is also apparent that progress against these challenges is being made. In IRA 5, a Husky robot platform equipped with the latest versions of all hardware and software attempted to address its environment semantically, recognizing building structures and objects, placing them in a world model, and then autonomously navigating with respect to the semantically populated world model in a manner specified by a tactical behavior specification language developed for semantic navigation. The overarching goal was for the robot to receive initial direction, much as a Soldier team member would be required to do, and then execute the direction autonomously, reacting to new environmental aspects at it progressed along a route consistent with the direction's intent. The distinction between this Soldier-like behavior and previous-generation navigation that relied principally on a priori maps, GPS, and local obstacle avoidance cannot be overstated.

IRA 5 did show some success toward the semantically driven goals. Fifty percent of the runs judged by human observation were at least partially successful in attaining the goal for the run, 85% at least partially met the navigation expectation, and 65% at least partially operated within defined constraints. Still, when all components are taken together, only 30% of the runs were

completely successful in all phases. This is likely how a human team member would evaluate the performance of the system. In addition to the obvious importance of the apparent result overall, it is also critical to report on the operation of the system components to better understand the reason for each success and root cause for each failure. To this end, evaluations were also conducted by the developers on the state of the world model throughout the run, followed by a conditional evaluation on navigation decisions subject to the state of that world model with errors in that model considered. Buildings, walls, and objects (vehicles and cones) had to be recognized and placed. In the case of buildings, parts of the building beyond sensor view also had to be predicted. The average number of vehicles seen per run was 5.6. Although tight controls of parking proximity were not in place, it is highly unlikely that there were this many vehicles within sensor view. The world model accurately reported the number of cones in 4 runs, underestimated in 6, and overestimated in 10. Underestimates might occur if the cone never comes within sensor view; for example, if the platform moves to the left of the building and there is a cone on the right. However, overestimates must be attributed to something else. Platform slippage, combined with an inability to recognize some previously seen objects, is one possible cause of over-counting cones, but this is not a satisfactory explanation for the over-counting of vehicles. Building predictions showed inconsistencies in 75% of the runs. Regarding navigation—conditional on the world model, 75% of the runs were successful, i.e., they were free of common errors detrimental to navigation, the most prevalent errors being grounding and prediction.

Adverse weather conditions and late integration, due in part to platform availability, adversely impacted the study. For example, the semantic labeling had not previously been trained in snow conditions, camera calibration methods were still being fine-tuned, and the major software components integrating the tactical behavior specifications, world model, and building prediction had been living on the same platform for too short a time. The after action review for IRA 5 led to several technology decisions to ready the system for the capstone study. Some were already in progress but not ready for IRA 5. A few of the changes are noted here to represent the program direction. An extended range for the LADAR is being implemented. This should help in all aspects of semantic navigation. A different sensor is being relied on for obstacle detection, relegating the original semantic labeling to longer range identifications of buildings and surroundings. A new open-space feature has been incorporated in building descriptions to improve building identification. Calibration procedures have been improved. Navigation solutions to slippage are being explored. A human-robot interface is being implemented so that it will be possible for the robot to receive and request clarifications on tactical behaviors. And perhaps most important, developers have had more opportunity in the months following IRA 5 to thoroughly test hardware and software in preparation for the RCTA capstone in late 2014.

5. References

1. Robotics collaborative technology alliance proposed FY12 annual program plan. Aberdeen (MD): Army Research Laboratory (US): 2012 Jan 18. [accessed 2014 Octobeer 27]. http://www.arl.army.mil/www/pages/392/RCTA_FY12_APP.pdf. See page 6 for a description of the hybrid cognitive/metric world model.

2. Dean R. Common world model for unmanned systems. Proceedings of SPIE 8741 Unmanned Systems Technology XV. Baltimore, MD; 2013 Apr 29. Bellingham (WA): The International Society for Optics and Photonics; 2013 May 17. [accessed 2015 Jan 15]. http://proceedings .spiedigitallibrary.org/proceedings.aspx?articleid=1689673.

3. Dean R, Oh J, Vinokurov J. Common world model for unmanned systems: phase 2. Proceedings of SPIE 9084 Unmanned Systems Technology XVI. Baltimore, MD; 2014 May 5; Bellingham (WA): The International Society for Optics and Photonics; 2014 Jun 24. [accessed 2015 Jan 15]. http://spie.org/Publications/Proceedings/Volume/9084.

4. Lennon C, Bodt B, Childers M, Camden R, Suppe A, Navarro-Serment L, Florea N. Performance evaluation of a semantic perception classifier. Aberdeen Proving Ground (MD): Army Research Laboratory (US); 2013 Sep. Report No.: ARL-TR-6653. Also available at: http://www.arl.army.mil/arlreports/2013/ARL-TR-6653.pdf.

5. Anderson J, Bothell D, Byrne M, Douglass S, Lebiere C, Yulin Q. An integrated theory of the mind. Psychological Review. 2004;111(4):1036–1060.

Appendix. Detailed Run Assessments

This appendix contains the scoring for 20 trials from Integrated Research Assessment 5, as shown in Figs. A-1 through A-66 and Tables A-1 through A-63. For building transformation images, such as Figure A-3, the colors are described in Table A-1. For images in which no transformation occurs, the colors are described in Table A-2. All scoring is in accordance with the guidelines set out in section 3 of the report. In the human-drawn record of the robot's path during the vignette, cones are depicted as circles, and vehicles intentionally placed in the scene are labeled with "veh." The 2 connected rectangles represent the church.

Table A-1 Legend for all images with building transformations

Object	Color and Line Type
Old building	Blue dashed line
Walls	Black
New building	Red solid line
Robot's previous path	Dark green solid line
Robot's planned path	Light green dashed line
Vehicles	Purple solid line
Cones	Orange solid line
Fire hydrants	Red solid line

Table A-2 Legend for all images without building transformations

Object	Color and line type
Building	Blue solid line
Walls	Black
Robot's previous path	Dark green solid line
Robot's planned path	Light green dashed line
Vehicles	Purple solid line
Cones	Orange solid line
Fire hydrants	Red solid line

A.1 Vignette 1, Path 1a, 16 December 2013 at 1420

Table A-3 Human observer scoring for vignette 1, path 1a, during the experiment

Aspect	Assessment	Comments
Appropriate navigation	Success	Command: "Navigate left of building to a cone near the building."
Following constraints	Success	Went left to the cone near the building.
Reaching goal	Success	Stopped within 1 m of cone C1.
Other	. . .	We do not know if it saw cone C2.
Overall	S	. . .

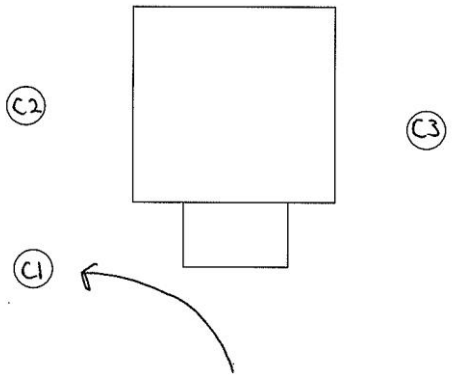

Fig. A-1 Layout for vignette 1 with observed path, path 1a, 16 December 2013 at 1420

Table A-4 Assessment of vignette 1, path 1a, based on world-model data

Assessed	Time	Code	Comments
Building	Final	R	Consistent with data and reasonable.
Nav. plan	Final	R	Appropriately follows constraints.
Vehicle	Final	. . .	4 recorded.
Cones	Final	. . .	1 recorded.

Analyst comments on vignette 1, path 1a: The robot perceived one cone and perceived the building to be in the same position throughout the run.

Table A-5 Assessment of vignette 1, path 1a, by researcher

Tactical Behavior Specification (TBS)	Navigate left of the building to a traffic cone near the building.
Result	Success
Explanation	The robot saw only one traffic cone that was near the building.

A.2 Vignette 1, Path 1b, 16 December 2013 at 1430

Table A-6 Human observer scoring for vignette 1, path 1b, during the experiment

Aspect	Assessment	Comments
Appropriate navigation	Success	Command: “Navigate left of building to a cone near the building.”
Following constraints	Success	Went left to the cone near the building.
Reaching goal	Success	Stopped within 1 m of C2.
Other	. . .	Cone C2 was moved after V1, path 1a, so the robot could see it.
Overall	S	. . .

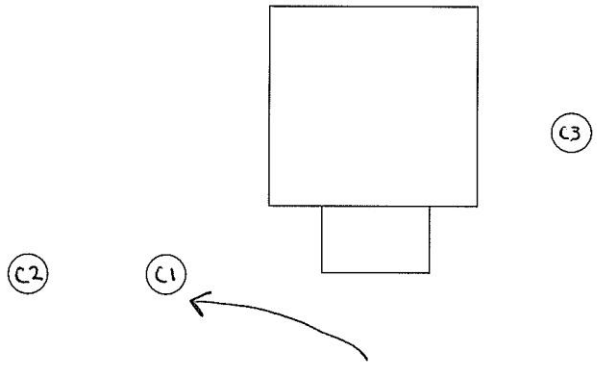

Fig. A-2 Layout for vignette 1, path 1b

Table A-7 Assessment of vignette 1, path 1b, based on world-model data

Assessed	Time	Code	Comments
Building	14:33:15.5	W, N	Fig. A-2
Nav. plan	Final	R	Appropriate planning
Cones	Final	. . .	6 cones observed, Fig. A-3
Vehicles	Final	. . .	6 vehicles observed, Fig. A-3

Analyst comments: Building prediction did not conform to wall observations, and there were false positives of cones and vehicles.

Table A-8 Assessment of vignette 1, path 1b, by researcher

TBS	Navigate left of the building to a traffic cone near the building.
Result	Success
Explanation	The robot saw multiple traffic cones and chose to go to the one near the building.

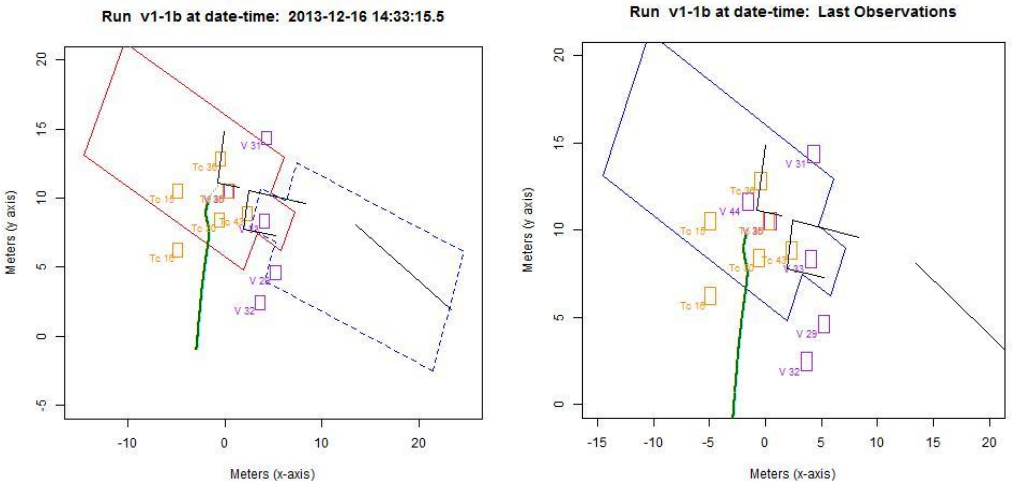

Fig. A-3 World-model reconstruction for vignette 1, path 1b

A.3 Vignette 1, Path 2-4, 16 December 2013 at 1435

Table A-9 Human observer scoring for vignette 1, path 2-4, during the experiment

Aspect	Assessment	Comments
Appropriate navigation	Partial success	Command: "Navigate right of building to a cone right of the building." Started right, turned left, then back right.
Following constraints	Success	. . .
Reaching goal	Success	Stopped within 1 m of C3.
Other	. . .	. . .
Overall	P 1	. . .

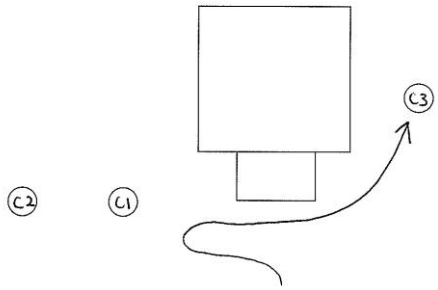

Fig. A-4 Layout for vignette 1, path 2-4

Table A-10 Assessments of vignette 1, path 2-4, based on world-model data

Assessed	Time	Code	Comments
Building	14:39:10.5	W, N	Fig. A-5
Vehicles	Final	. . .	12 observed, Fig. A-6 (not all shown)
Cones	Final	. . .	8 observed, Fig. A-6
Nav. plan	14:38:03.8	R	It is hard to determine which is the right side of the building, but the plan is appropriate when the building is correctly predicted, Fig. A-7.

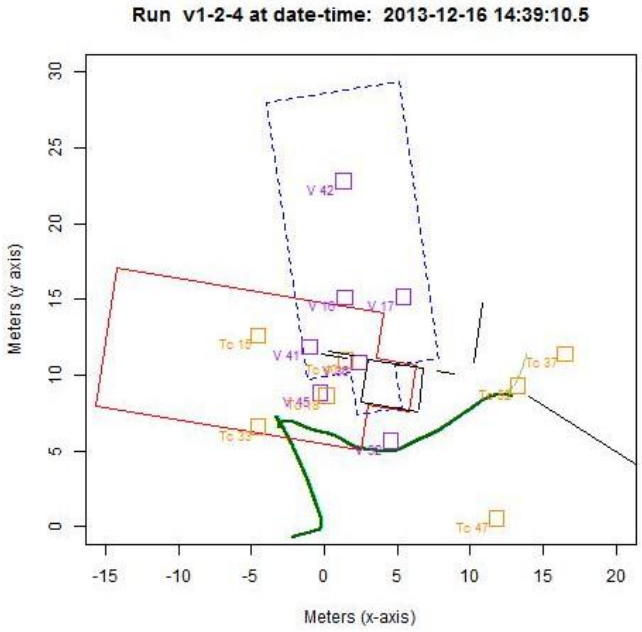

Fig. A-5 Building transformation from vignette 1, path 2-4

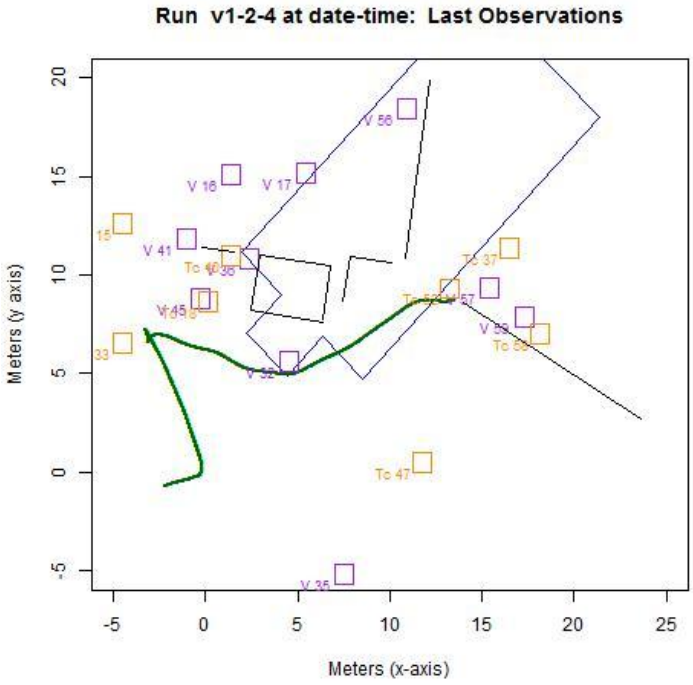

Fig. A-6 Final model from vignette 1, path 2-4

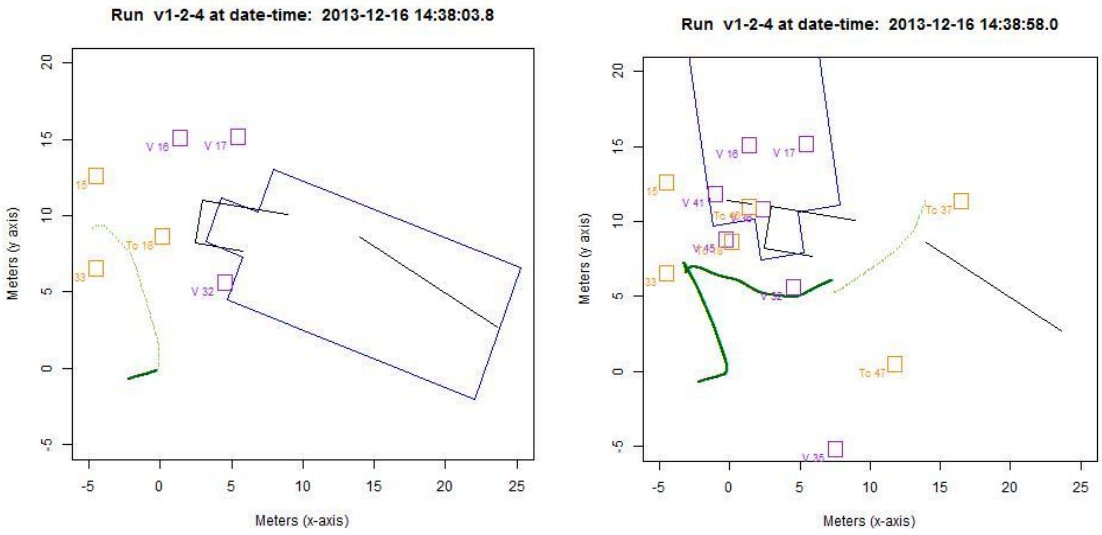

Fig. A-7 Navigation from vignette 1, path 2-4

Analyst comments: The building prediction was inconsistent with wall positions. The navigation planner has a delay of several seconds between building re-prediction and route re-planning.

Table A-11 Assessment of vignette 1, path 2-4, by researcher

TBS	Navigate right of the building to the traffic cone right of the building.
Result	Success
Explanation	[GR] Due to a grounding error in the beginning, the robot navigated to the left first, trying to go to a traffic cone that it could see from the front. As the robot drove near the building, it saw the other traffic cone on the right, re-plans, and finishes at the correct goal location. The first part of the behavior can be considered a constraint violation.

A.4 Vignette 1, Path 2-4b, 16 December 2013 at 1443

Table A-12 Human observer scoring for vignette 1, path 2-4b, during the experiment

Aspect	Assessment	Comments
Appropriate navigation	Partial success	Command: "Navigate right of building to a cone right of the building."
Following constraints	Fail	Went directly to cone near the building on the left.
Reaching goal	Fail	Stopped within 1 m of cone 1 (wrong cone).
Other	. . .	. . .
Overall	F 3	. . .

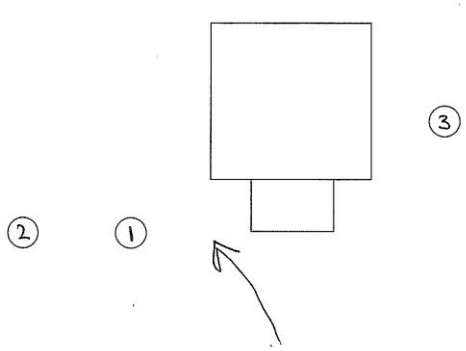

Fig. A-8 Layout for vignette 1, path 2-4b,
16 December 2013 at 1443

Table A-13 Assessment of vignette 1, path 2-4b, based on world-model data

Assessed	Time	Code	Comments
Building	14:44:20.7	R	Although incorrect, this counts as being consistent with wall position, Fig. A-9.
Vehicles	Final	. . .	4 observed (not all shown), Fig. A-10.
Cones	Final	. . .	5 observed, Fig. A-10.
Nav. plan	14:43:51.4	C	Fig. A-11

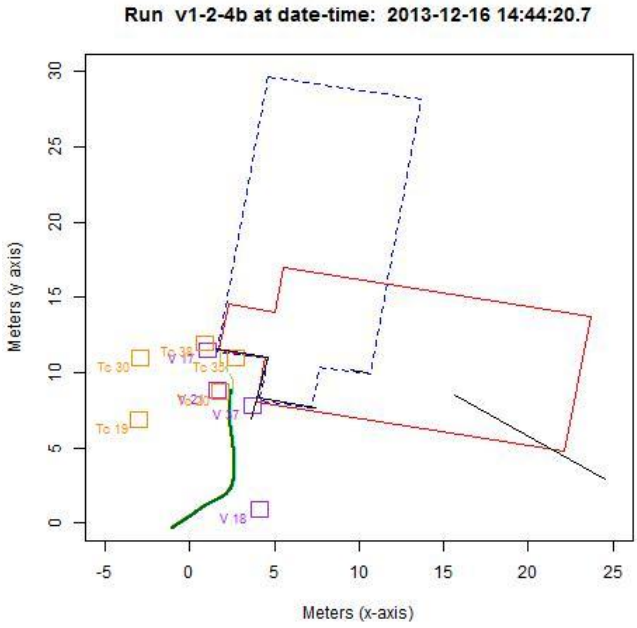

Fig. A-9 Building transformation from vignette 1,
path 2-4b

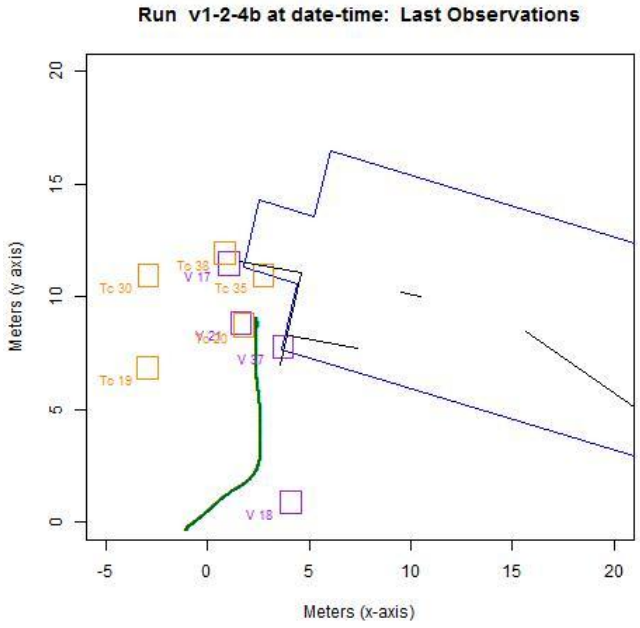

Fig. A-10 Final model from vignette 1, path 2-4b

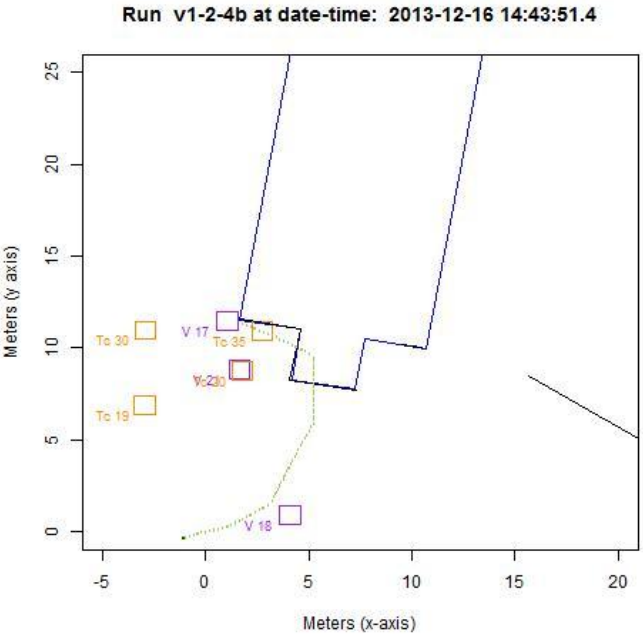

Fig. A-11 Navigation planning from vignette 1, path 2-4b

Analyst comments: The building predictor was consistent with wall observations and thus should be considered successful in this run. The robot seemed to ignore the constraint about the cone being right of the building. It did not make it far enough to see the correct cone.

Table A-14 Assessment of vignette 1, path 1-2b, by researcher

TBS	Navigate right of the building to the traffic cone right of the building
Result	Fail
Explanation	[GR] Grounding error

A.5 Vignette 1, Path 3-5, 16 December 2013 at 1455

Table A-15 Human observer scoring for vignette 1, path 3-5, during the experiment

Aspect	Assessment	Comments
Appropriate navigation	Partial success	Command: "Navigate left of building to a cone left of the building."
Following constraints	Fail	Went left, then back right, stopped in front of steps.
Reaching goal	Fail	Stopped far from cone.
Other	. . .	Cone 2 was moved back to the left of the building.
Overall	F 3	. . .

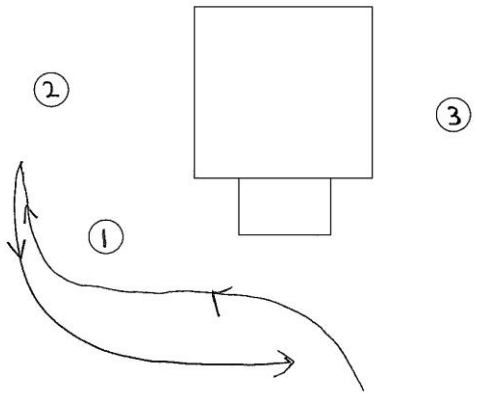

Fig. A-12 Layout for vignette 1, path 3-5

Table A-16 Assessment of vignette 1, path 3-5, based on world-model data

Assessed	Time	Code	Comments
Building	14:54:55.5	N	It could have kept better to the walls, but the building was aligned with two of them, Fig. A-13.
Cones	Final	. . .	4 observed, Fig. A-14.
Vehicles	Final	. . .	4 observed, Fig. A-14.
Nav. plan	14:55:18.3	O	The change of course could have resulted from the changed building position; the robot might have been heading through the building to the right side of the building, Fig. A-15.

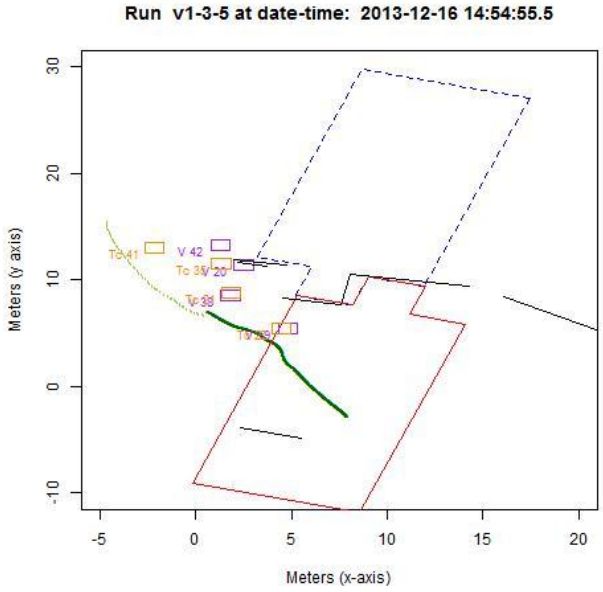

Fig. A-13 Building transformation from vignette 1, path 3-5

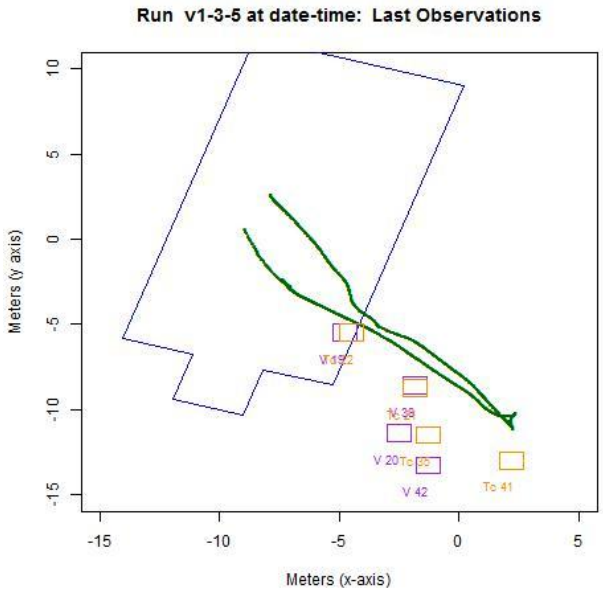

Fig. A-14 Final model for vignette 1, path 3-5

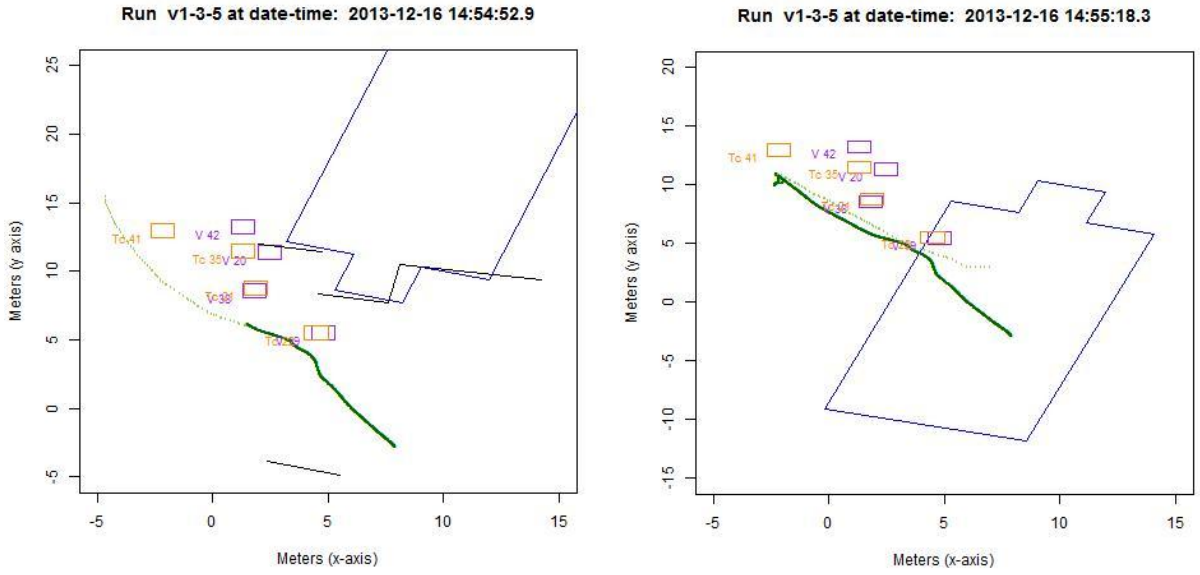

Fig. A-15 Navigation planning from vignette 1, path 3-5

Analyst comments: The building prediction followed wall positions well, and navigation planning performed is difficult to judge given building prediction. The building prediction placed the building in a position in which the robot had already traveled. Some level of reasoning about building position would have been useful here.

Table A-17 Assessment of vignette 1, path 3-5, by researcher

TBS	Navigate left of the building to the traffic cone left of the building
Result	Fail
Explanation	• The robot saw only the cones that were placed on the left front corner of the building and aimed at one of them. • The robot saw the (correct) traffic cone left of the building, re-planned, and reached the goal. • [GR] A grounding error occurred, and the robot ended up finishing at an incorrect traffic cone at left front of the building.

A.6 Vignette 1, Path 3-5b, 16 December 2013 at 1458

Table A-18 Human observer scoring for vignette 1, path 3-5b, during the experiment

Aspect	Assessment	Comments
Appropriate navigation	Partial success	Command: "Navigate left of building to a cone left of the building."
Following constraints	Fail	Started left, then right, then left toward fire hydrant, then stuck in snow
Reaching goal	Fail	Stuck in the snow
Other	. . .	Started to the right of its previous position
Overall	F 3	. . .

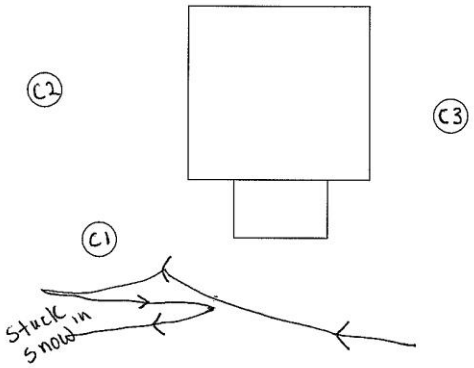

Fig. A-16 Layout for vignette 1, path 3-5b,
16 December 2013 at 1458

Table A-19 Assessment of vignette 1, path 3-5b, based on world-model data

Assessed	Time	Code	Comments
Building	15:02:49.9	W, N	Fig. A-18
Nav. plan	15:01:50.3	O	When it sees the cone, it keeps heading there, ignoring the building predicted there for 12 s, Fig. A-17.
Cones	Final	. . .	6 observed, Fig. A-18
Vehicles	Final	. . .	12 observed, Fig. A-18
Fire hydrant	Final	. . .	1 observed, Fig. A-18

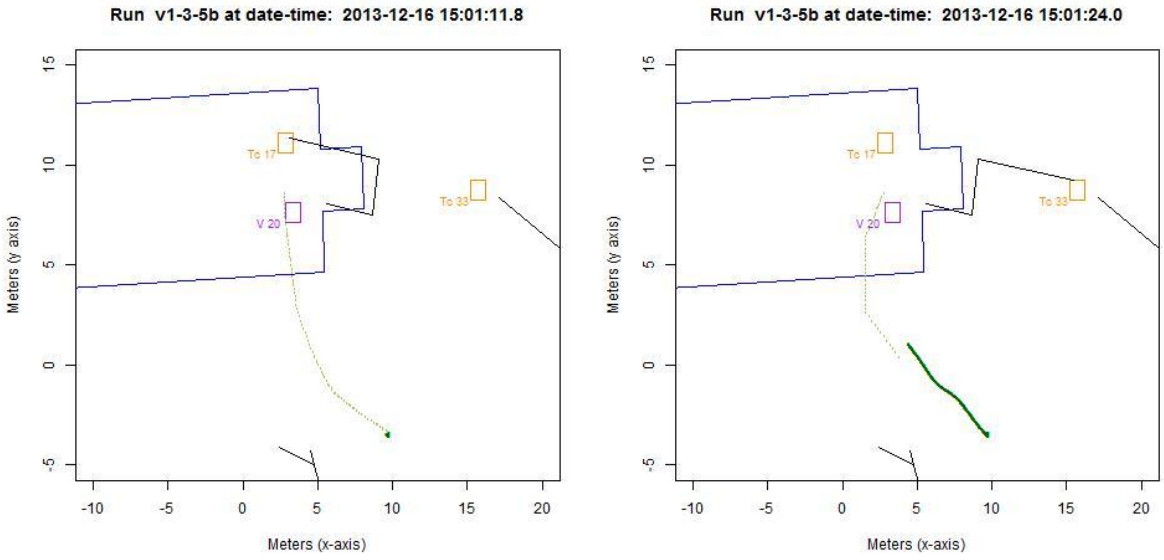

Fig. A-17 Navigation planning from vignette 1, path 3-5b

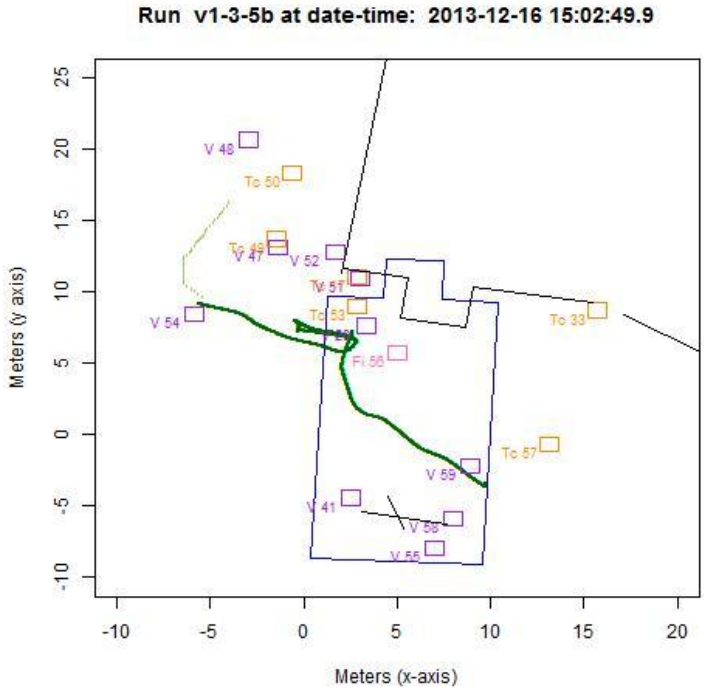

Fig. A-18 Final model from vignette 1, path 3-5b

Analyst comments: The building prediction did not match well with the observed wall positions. The prediction seems to have been ignored by the path planner when it could see the cone. It is possible that the multiple false positives on objects could come, in part, from the robot's difficulty with localization. This could have resulted from the skidding and spinning in the snow, observed by the experimenters, but this is uncertain given that the robot did a generally good job of tracking where it was, plus there were many false positives on other runs as well in which no such loss of control occurred.

Table A-20 Assessment of vignette 1, path 3-5b, by researcher

TBS	Navigate left of the building to the traffic cone left of the building.
Result	Fail
Explanation	• The robot sees only the cones that are placed on the left front corner of the building and plans for that goal. • [EN] The robot sees the (correct) traffic cone left of the building, re-plans but fails to reach the goal due to snow piles on the path. • [GR] Grounding error occurs, and the robot re-plans to go to the other traffic cone left front of the building.

A.7 Vignette 1, Path 6, 16 December 2013 at 1506

Table A-21 Human observer scoring for vignette 1, path 6, during the experiment

Aspect	Assessment	Comments
Appropriate navigation	Success	Command: "Navigate left of building to a cone near the building."
Following constraints	Success	. . .
Reaching goal	Partial success	Stopped 2–3 m from cone.
Other	. . .	Started to the right of its previous position, and cone C2 was moved farther from the building.
Overall	P 1	. . .

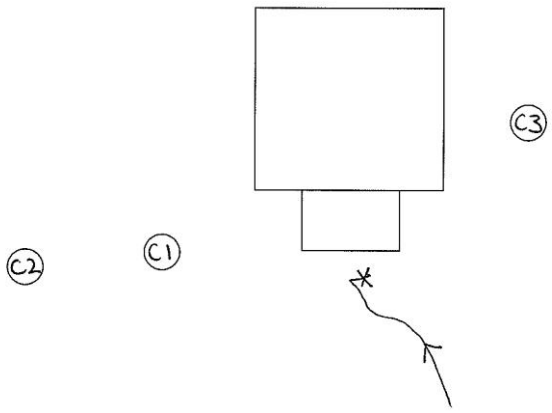

Fig. A-19 Layout for vignette 1, path 6

Table A-22 Assessment of vignette 1, path 6, based on world-model data

Assessed	Time	Code	Comments
Building	15:10:25.0	W, N	Fig. A-20
Building	15:10:43.5	W, N	Fig. A-20
Nav. plan	15:10:30.0 15:10:18.4	R	Appropriately planned paths as the building moved, Figs. A-21 and A-22
Cones	Final	. . .	3 observed
Vehicles	Final	. . .	3 observed

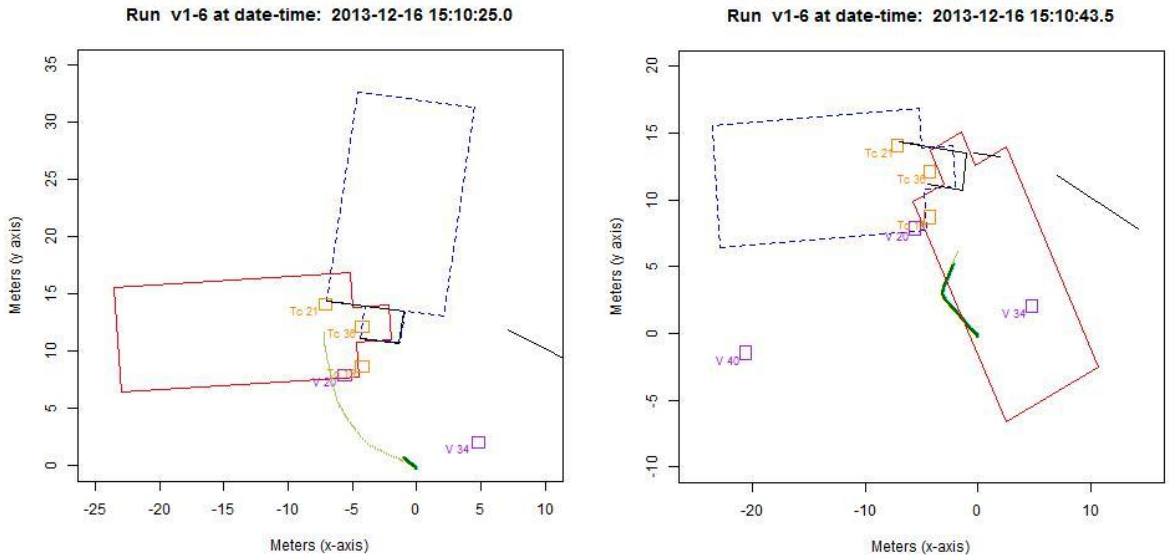

Fig. A-20 Building transformations from vignette 1, path 6

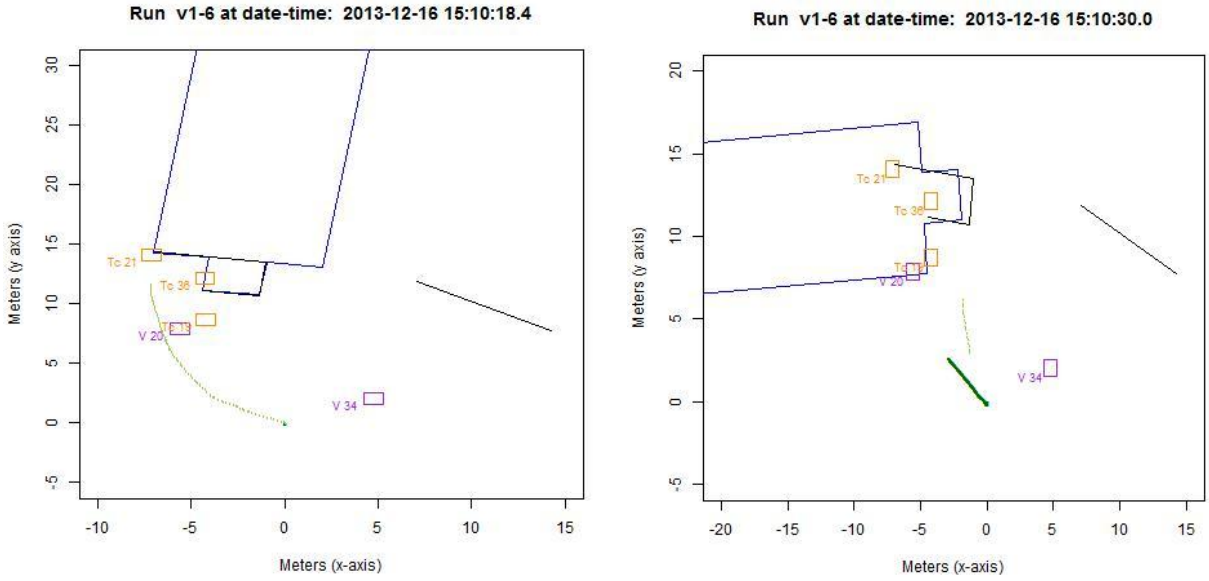

Fig. A-21 Navigation planning from vignette 1, path 6

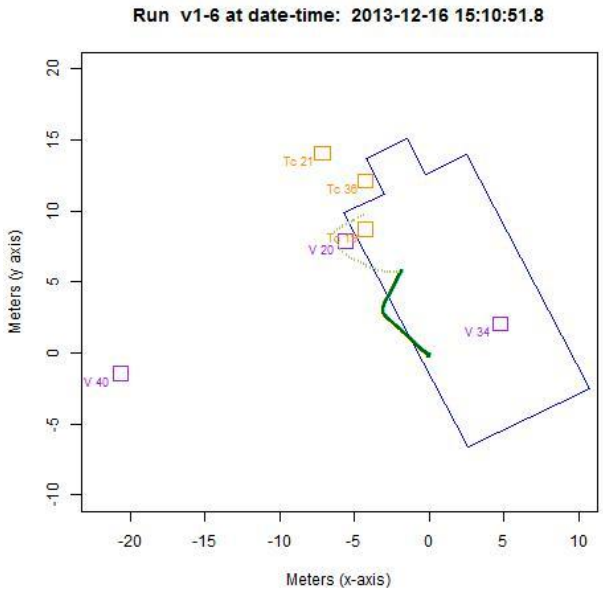

Fig. A-22 Final model from vignette 1, path 6

Analyst comments: The navigation planning was appropriate given the building position. The building predictors differed from wall observations and positioned the building on spots on which the robot had traveled.

Table A-23 Assessment of vignette 1, path 6, by researcher

TBS	Navigate left of the building to a traffic cone near the building.
Result	Success
Explanation	Desired behavior was observed. [OD] Due to fusion errors, however, the robot saw a few false positives near the true-positive traffic cone.

A.8 Vignette 1, Path 6b, 16 December 2013 at 1517

Table A-24 Human observer scoring for vignette 1, path 6b, during the experiment

Aspect	Assessment	Comments
Appropriate navigation	Fail	Command: “Navigate left of building to a cone near the building.”
Following constraints	Success	Started going up the church steps, and was stopped to avoid damage.
Reaching goal	Fail	. . .
Other	. . .	. . .
Overall	F 3	. . .

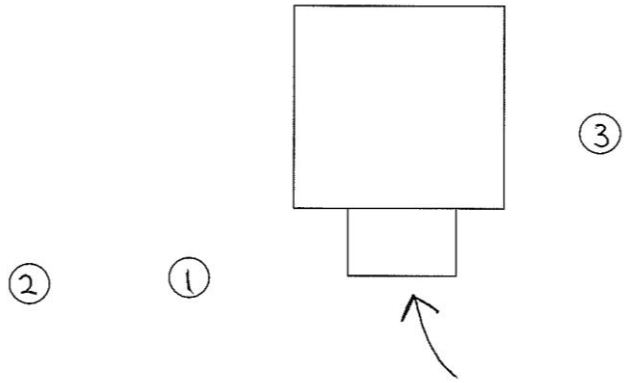

Fig. A-23 Layout for vignette 1, path 6b, 16 December 2013 at 1517

Table A-25 Assessment of vignette 1, path 6b, based on world-model data

Assessed	Time	Code	Comments
Building	15:18:57.7	W	Fig. A-24
Building	15:19:27.3	W, N	Fig. A-24
Cones	Final	. . .	5 observed, Fig. A-25
Vehicles	Final	. . .	3 observed, Fig. A-25

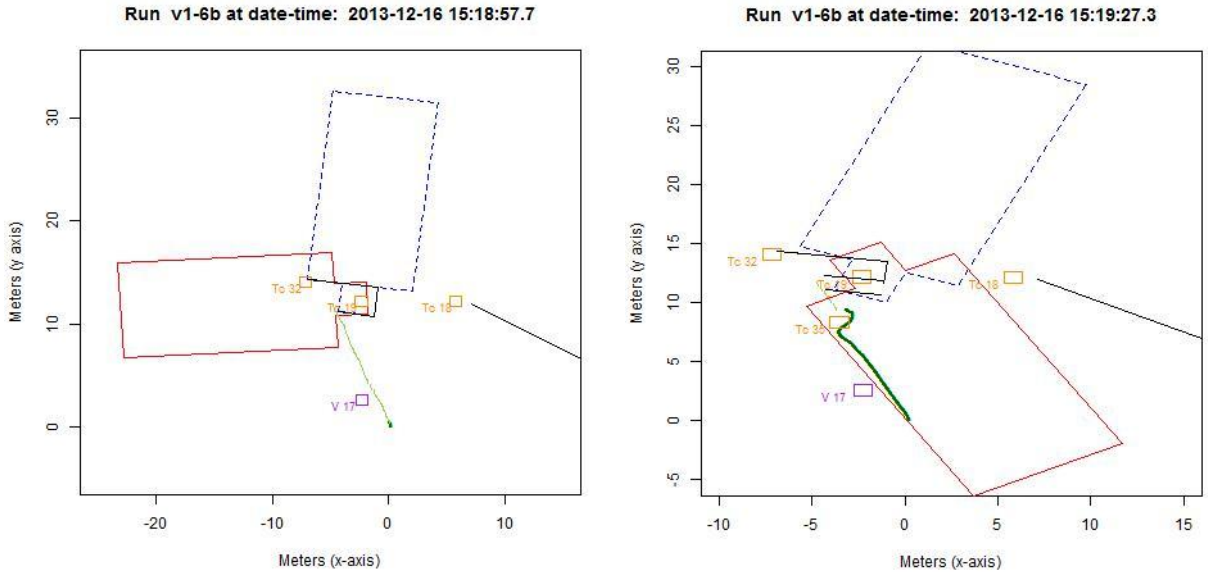

Fig. A-24 Building transformations from vignette 1, path 6b

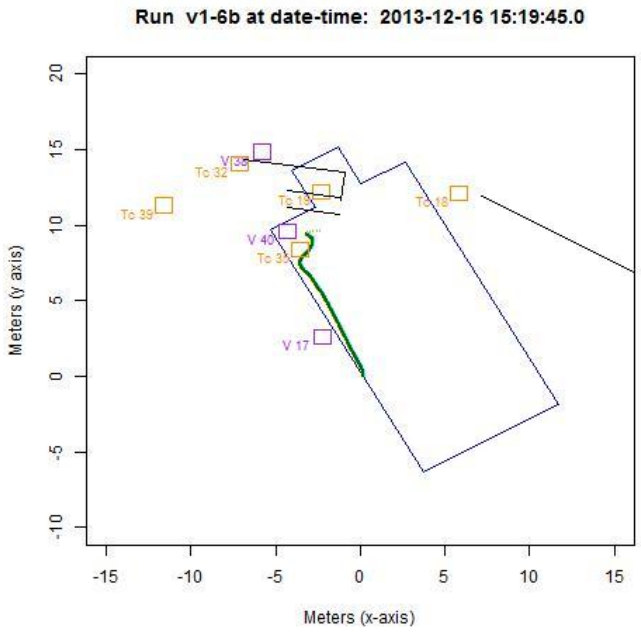

Fig. A-25 Final model from vignette 1, path 6b

Analyst comments: Building prediction had difficulty aligning with observed walls. The navigation planner made reasonable decisions given what was in the world-model. The robot failed to avoid the stairs of the church.

Table A-26 Assessment of vignette 1, path 6b, by researcher

TBS	Navigate left of the building to a traffic cone near the building.
Result	Partial success
Explanation	[PL] The correct goal was identified but the robot was stuck on the front steps of the building at midway.

A.9 Vignette 1, Path 6c, 16 December 2013 at 1525

Table A-27 Human observer scoring for vignette 1, path 6c, during the experiment

Aspect	Assessment	Comments
Appropriate navigation	Success	Command: "Navigate left of building to a cone near the building."
Following constraints	Success	. . .
Reaching goal	Success	Navigated to within 1 m of the correct cone.
Other	. . .	. . .
Overall	S	. . .

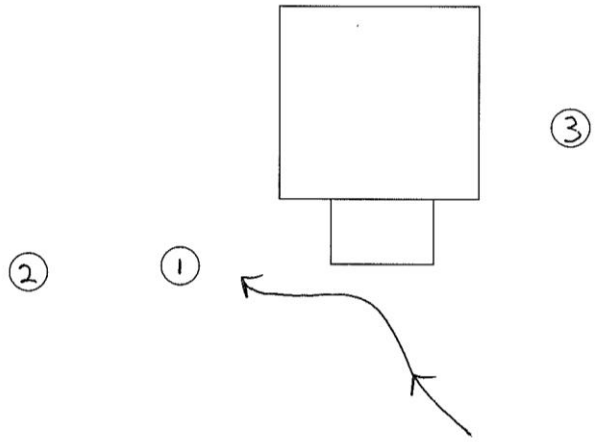

Fig. A-26 Layout for vignette 1, path 6c,
16 December 2013 at 1525

Table A-28 Assessment of vignette 1, path 6c, based on world-model data

Assessed	Time	Code	Comments
Building	Final	R	Prediction starts correct and does not change much during the run, Fig. A-27
Nav. plan	Final	R	Appropriate throughout the run, Fig. A-27
Cones	Final	. . .	4 observed, Fig. A-27
Vehicles	Final	. . .	7 observed, Fig. A-27

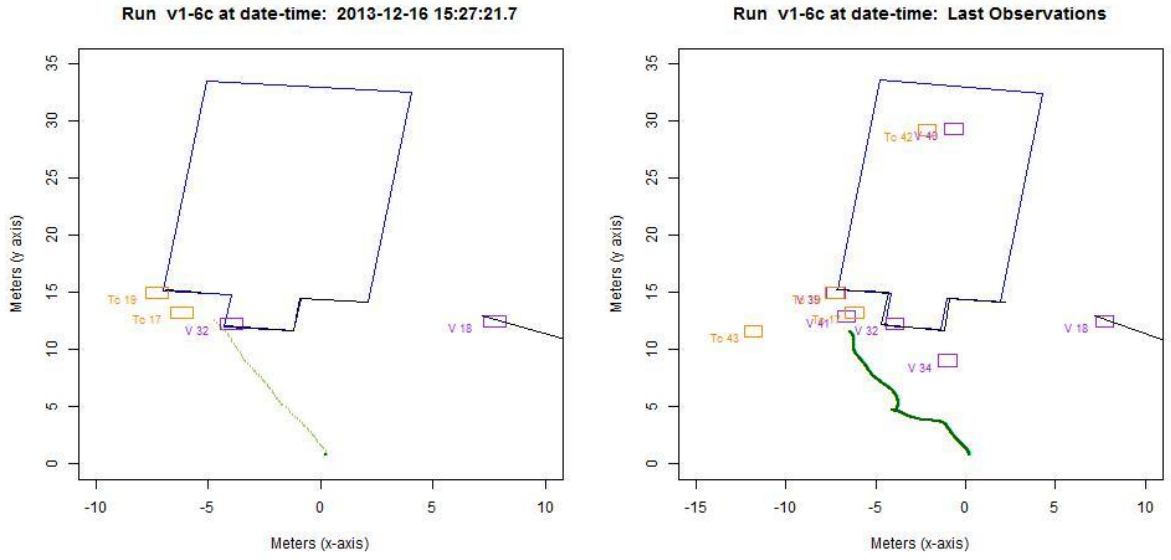

Fig. A-27 World-model reconstruction for vignette 1, path 6c

Analyst comments: The robot only saw cone 1 throughout the run and had reached the end before it saw cone 2.

Table A-29 Assessment of vignette 1, path 6c, by researcher

TBS	Navigate left of the building to a traffic cone near the building.
Result	Success
Explanation	The robot detected a traffic cone near the building when the command was given. The shape of the path generated satisfied the "left of the building" constraint. [OD] There were a few false positive cars.

A.10 Vignette 4, Path 1a, 16 December 2013 at 1535

Table A-30 Human observer scoring for vignette 4, path 1a, during the experiment

Aspect	Assessment	Comments
Appropriate navigation	Fail	Command: "Navigate left of building to a cone behind the building."
Following constraints	Fail	Started left, then saw a cone to the right of the building and went toward it.
Reaching goal	Fail	Stopped by operators.
Other	. . .	. . .
Overall	F 3	. . .

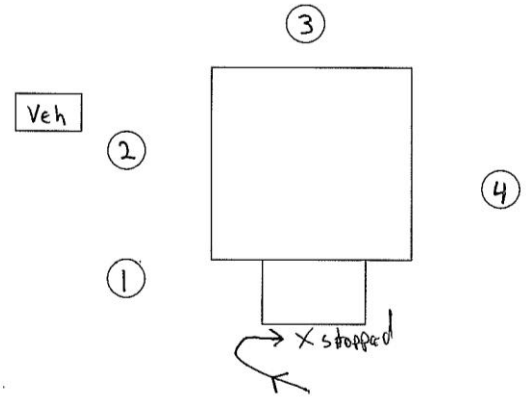

Fig. A-28 Layout for vignette 4, path 1a

Table A-31 Assessment of vignette 4, path 1a, based on world-model data

Assessed Cat	Time	Code	Comments
Building	15:36:35.0 15:35:47.2	W	Fig. A-29
Nav. plan	15:36:29.0	C	Change from left to right, Fig. A-30
Cones	Final	. . .	2 observed , Fig. A-31
Vehicles	Final	. . .	8 observed, Fig. A-31

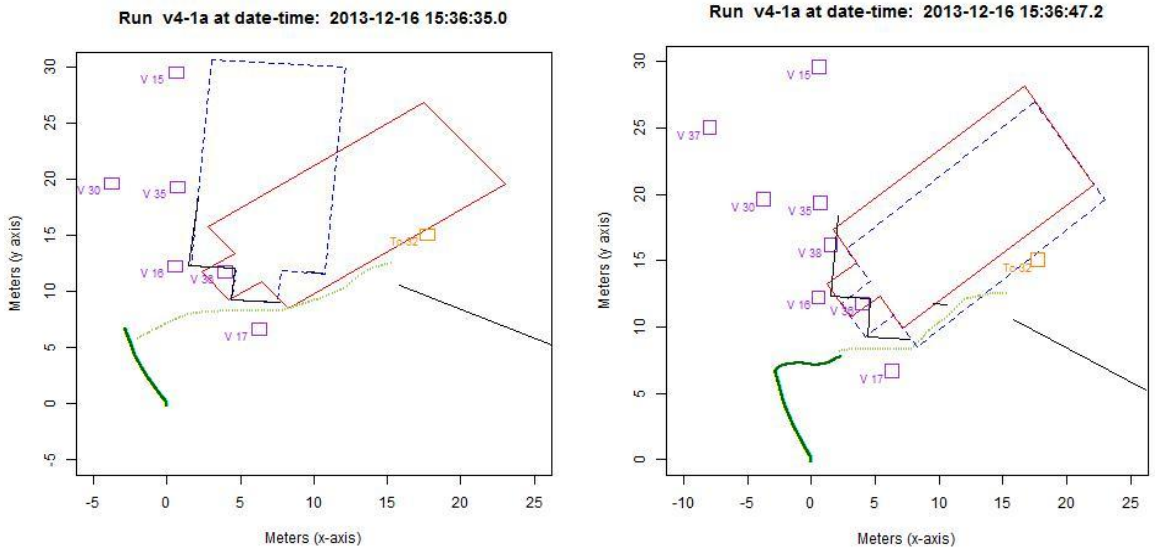

Fig. A-29 Building transformations from vignette 4, path 1a

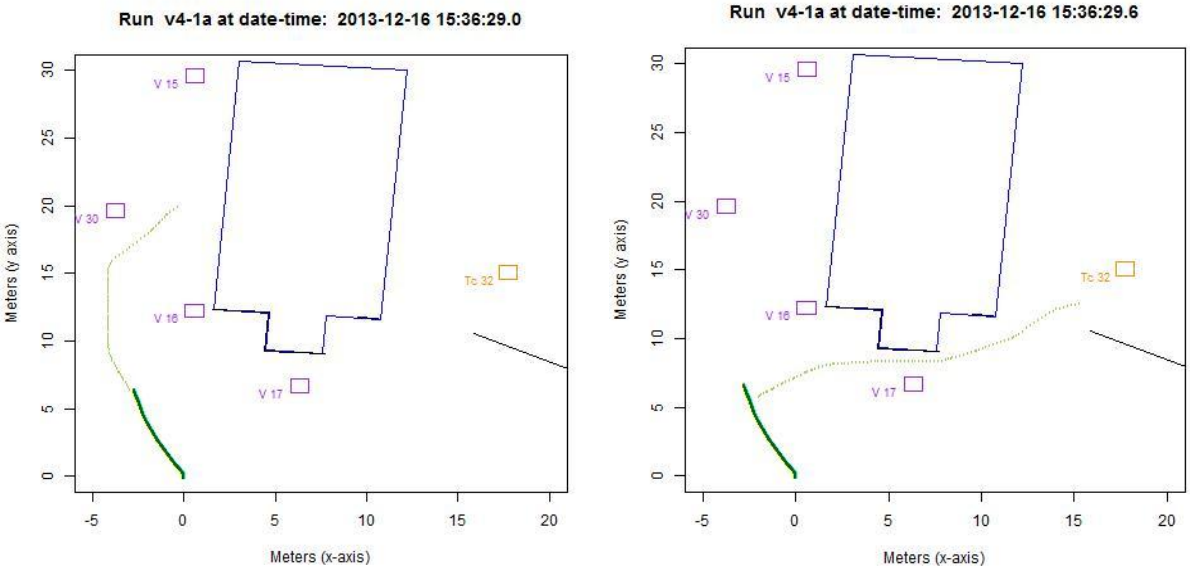

Fig. A-30 Navigation planning from vignette 4, path 1a

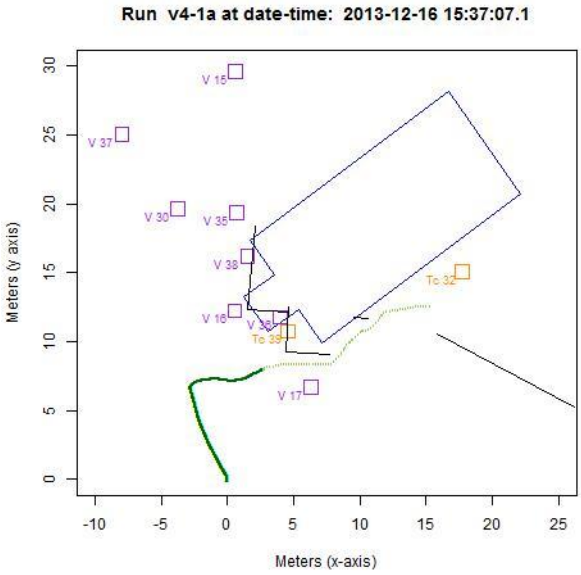

Fig. A-31 Final model for vignette 4, path 1a

Analyst comments: Although the building prediction did deviate from the observed walls, the navigation planner violated the left side constraint before that occurred. Again, there are a lot of false positive vehicles.

Table A-32 Assessment of vignette 4, path 1a, by researcher

TBS	Navigate left of the building to a traffic cone behind the building.
Result	Fail
Explanation	[GR] The robot could not see any traffic cones from its starting position, so it started navigating to the left of the building. When the robot saw a traffic cone on the right, it incorrectly grounded that as a goal.

A.11 Vignette 4, Path 1b, 17 December 2013 at 1330

Table A-33 Human observer scoring for vignette 4, path 1b, during the experiment

Aspect	Assessment	Comments
Appropriate navigation	Success	Command: "Navigate left of building to a cone behind the building."
Following constraints	Success	. . .
Reaching goal	Partial success	Got stuck in the snow 3–5 m from the goal.
Other	. . .	After it got stuck, we shoveled snow off the hill it got stuck on.
Overall	P 1	. . .

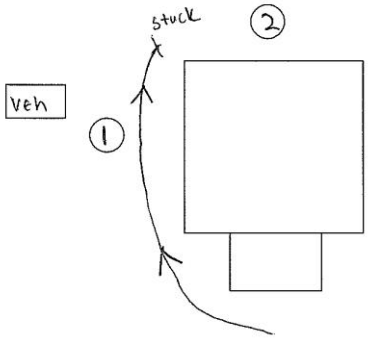

Fig. A-32 Layout for vignette 4, path 1b

Table A-34 Assessment of vignette 4, path 1b, based on world-model data

Assessed	Time	Code	Comments
Building	13:31:52.8	R	A building shift that is not inconsistent with walls, but could use them better, Fig. A-33.
Nav. plan	Final	R	. . .
Localization	13:33:27.0	O	The robot was stuck at the corner when the run finished. Fig. A-34 might depict localization failure, or the robot might have been pushed free and continued collecting data. This is a possible failure.
Cones	Final	. . .	2 observed but did not see the cone in back of the building, Fig. A-34.
Vehicles	Final	. . .	2 observed but missed the real one, Fig. A-34.

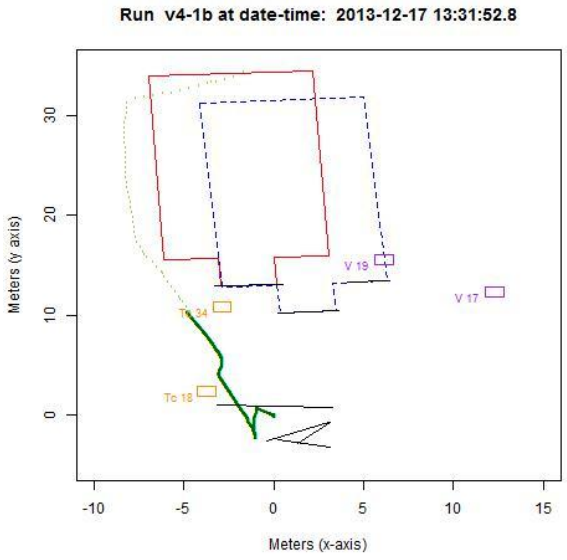

Fig. A-33 Building transformation from vignette 4, path 1b

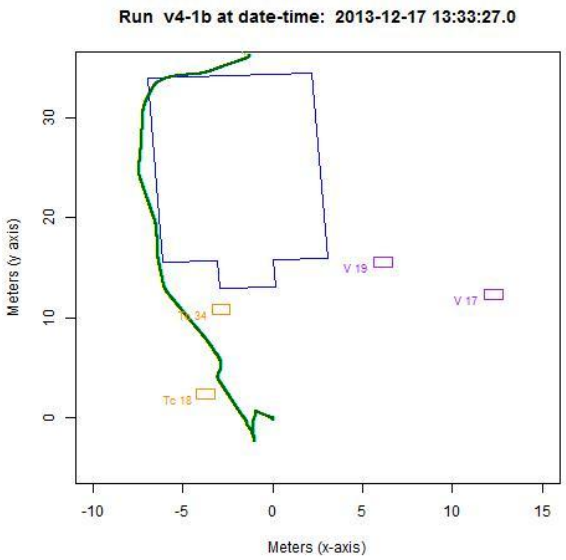

Fig. A-34 Final model from vignette 4, path 1b

Analyst comments: Due to the snow and ice, the robot eventually got stuck. It may be that difficulties with localization are causing problems for the building predictor but those do not seem to have caused a problem in this run. Given localization difficulties, the building shift in Fig. A-33 may have been a correction for position and is not counted as an error.

Table A-35 Assessment of vignette 4, path 1b, by researcher

TBS	Navigate left of building to a cone behind the building.
Result	Partial success.
Explanation	[OD] The robot's initial goal was an unknown object predicted behind the building because it could not see the cone behind the building from its start position. The robot failed to detect the traffic cone behind the building and finished at the predicted unknown object.

A.12 Vignette 4, Path 1c, 17 December 2013 at 1346

Table A-36 Human observer scoring for vignette 4, path 1c, during the experiment

Aspect	Assessment	Comments
Appropriate navigation	Success	Command: “Navigate left of building to a cone behind the building.”
Following constraints	Success	. . .
Reaching goal	Success	Within 1 m.
Other	. . .	. . .
Overall	S	. . .

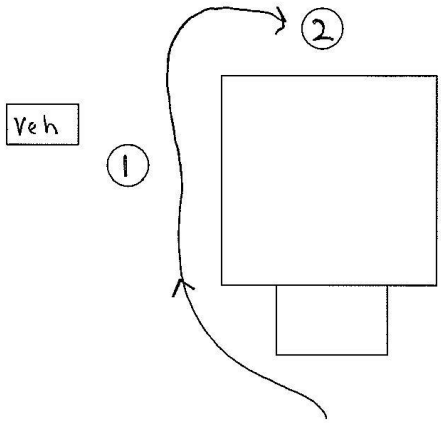

Fig. A-35 Layout for vignette 4, path 1c

Table A-37 Assessment of vignette 4, path 1c, based on world-model data

Assessed	Time	Code	Comments
Building	13:48:46.4 13:49:14.7	W, N	Fig. A-36
Nav. plan	13:48:57.3	O, R	The planner calmly ignores the building prediction to continue on its path, Fig. A-37.
Cones	Final	. . .	2 observed and one did not exist, Fig. A-37.
Vehicles	Final	. . .	0 observed, Fig. A-37.
Fire hydrant	Final	. . .	2 observed, neither existed, Fig. A-37.

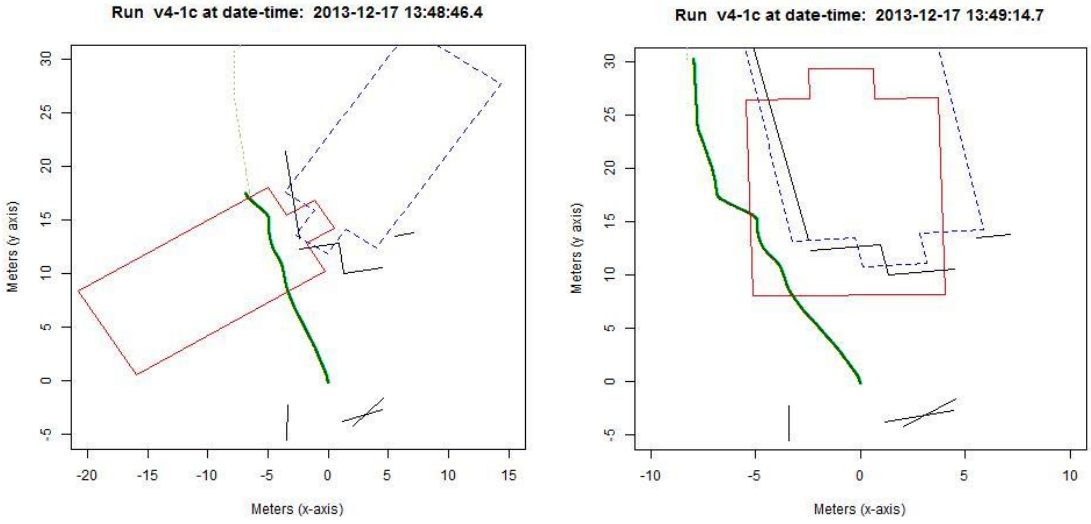

Fig. A-36 Building transformations from vignette 4, path 1c

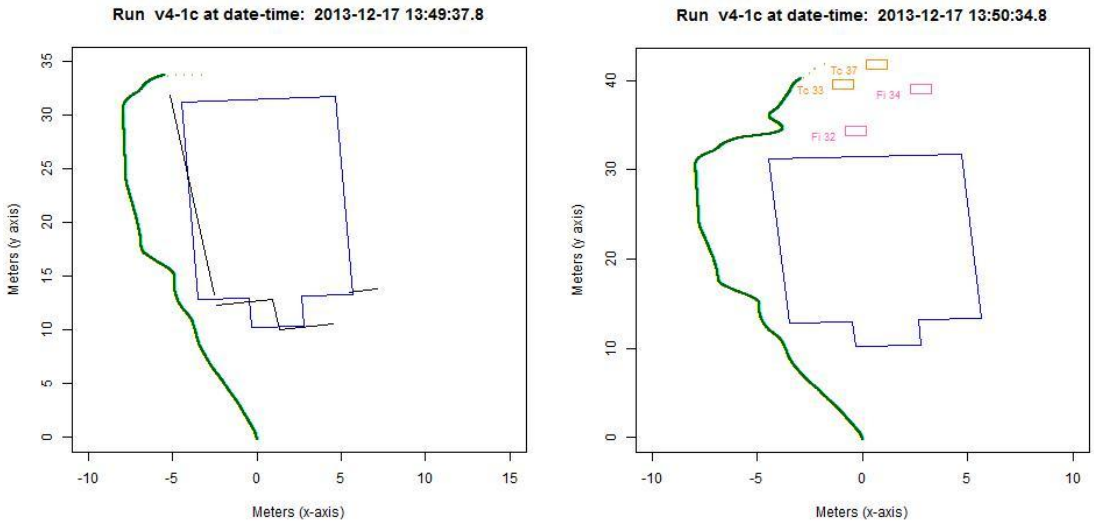

Fig. A-37 Navigation planning from vignette 4, path 1c

Analyst comments: The navigation planner ignored many of the building repositions and may have a waiting time of a few seconds between building repositioning and route re-planning. This delay seems to work well, as the speeds are low and building prediction is inconsistent.

Table A-38 Assessment of vignette 4, path 1c, by researcher

TBS	Navigate left of building to a traffic cone behind the building.
Result	Success
Explanation	The robot used an unknown object predicted behind the building as an initial goal. When the robot was behind the building, it detected a traffic cone behind the building, re-planned, and finished at the correct goal. [OD] The robot failed to detect a car and another traffic cone on the far left of the building. [GR] For brief period near frame 400, the grounder selected another unknown object on the right as a goal.

A-13 Vignette 4, Path 2, 17 December 2013 at 1400

Table A-39 Human observer scoring for vignette 4, path 2, during the experiment

Aspect	Assessment	Comments
Appropriate navigation	Partial success	Command: "Navigate left of building to a cone behind the building, covertly."
Following constraints	Partial success	. . .
Reaching goal	Fail	. . .
Other	. . .	Started out right, then spun in the snow.
Overall	F 1	. . .

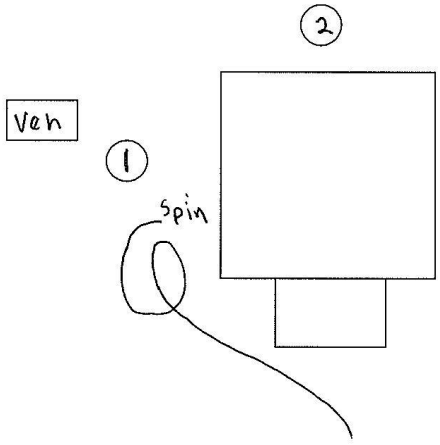

Fig. A-38 Layout for vignette 4, path 2

Table A-40 Assessment of vignette 4, path 2, based on world-model data

Assessed	Time	Code	Comments
Building	14:01:52.2	W, N	Fig. A-39
Nav. plan	14:01:59.0	R	It plans navigation to its own left around the building, Fig. A-40.
Cones	Final	. . .	2 observed , Fig. A-41.
Vehicles	Final	. . .	10 observed (not all shown), Fig. A-41.
Navigation	Final	O	The robot spins in circles while knowing that it is spinning in circles, Fig. A-41.

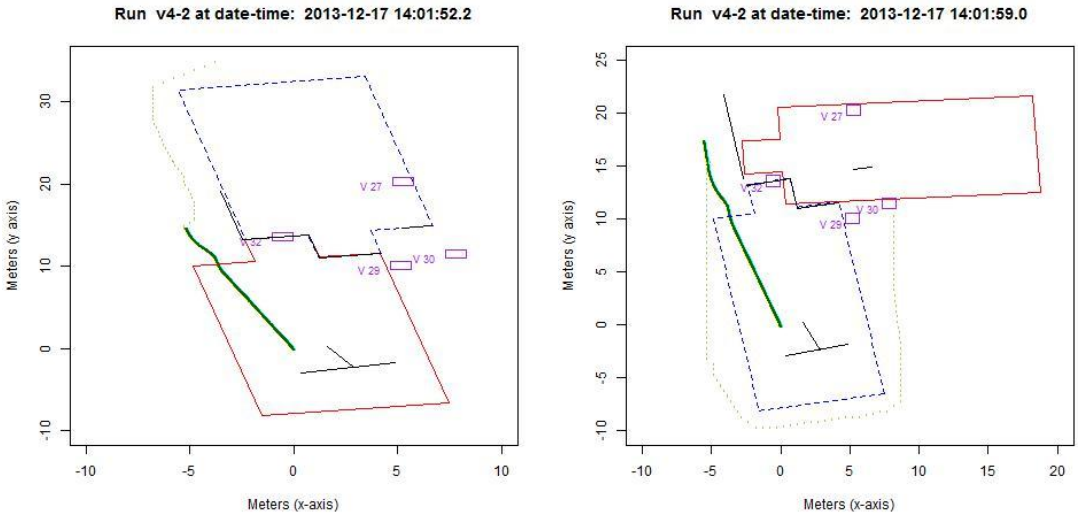

Fig. A-39 Building transformations from vignette 4, path 2

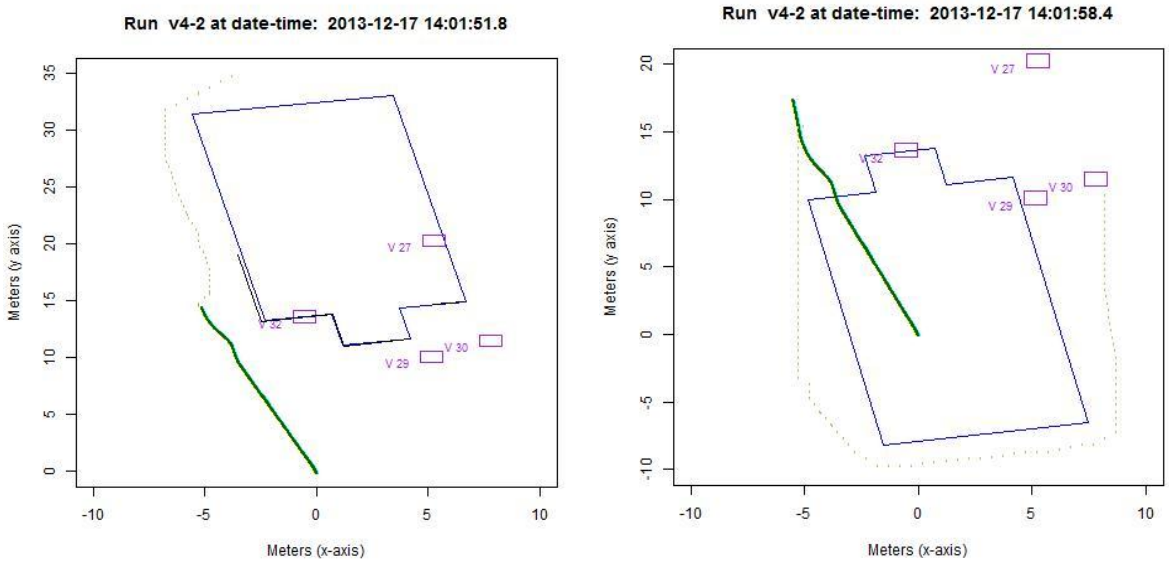

Fig. A-40 Navigation planning from vignette 4, path 2

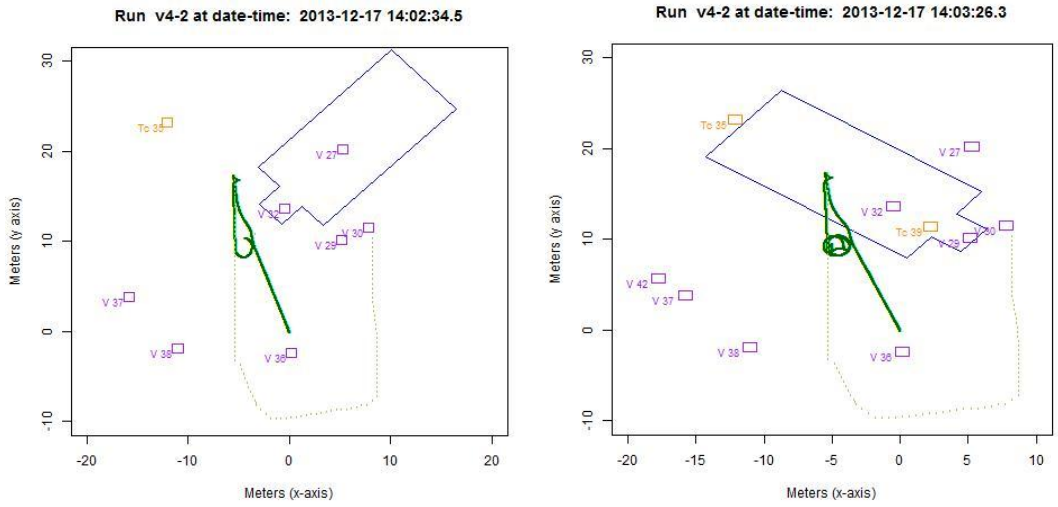

Fig. A-41 World-model reconstruction for vignette 4, path 2

Analyst comments: The navigation planner set up a good route but the robot kept traveling in circles. It correctly tracked its movement during this time but did not alter its behavior.

Table A-41 Assessment of vignette 4, path 2, by researcher

TBS	Navigate covertly left of building to a cone behind the building.
Result	Fail
Explanation	[PR] The building was incorrectly predicted on the wrong side.

A-14 vignette 4, path 3, 17 December 2013 at 1420

Table A-42 Human observer scoring for vignette 4, path 3, during experiment

Aspect	Assessment	Comments
Appropriate navigation	Partial success	Command: “Navigate left of building to a cone near vehicle.”
Following constraints	Fail	. . .
Reaching goal	Fail	. . .
Other	. . .	. . .
Overall	F 3	. . .

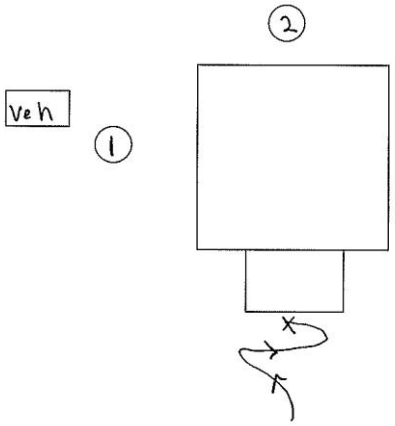

Fig. A-42 Layout for vignette 4, path 3, 17 December 2013 at 1420

Table A-43 Assessment of vignette 4, path 3, based on world-model data

Assessed	Time	Code	Comments
Building	14:22:38.1	W	Fig. A-43
Nav. plan	14:22:16.3 14:23:33.3	C	The planner changes direction to violate the constraint, Figs. A-44 and A-45.
Cones	Final	. . .	0 observed, Fig. A-45.
Vehicles	Final	. . .	3 observed but not the one near the cone, Fig. A-45.

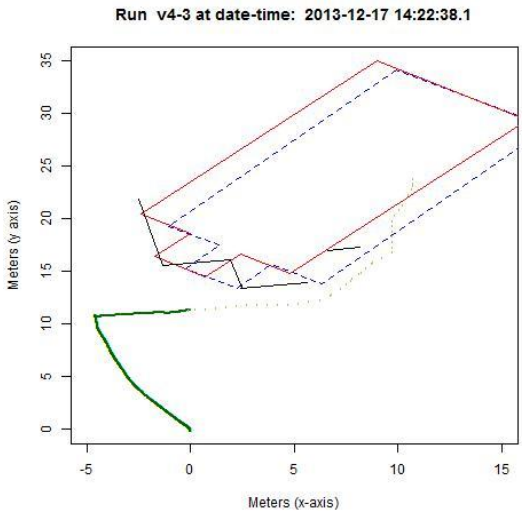

Fig. A-43 Building transformation from vignette 4, path 3

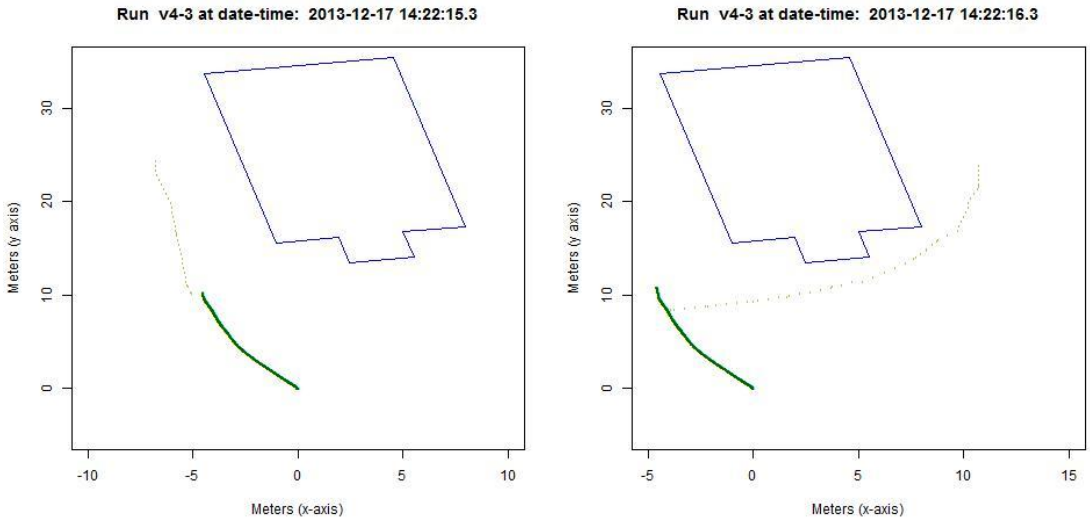

Fig. A-44 Navigation planning from vignette 4, path 3

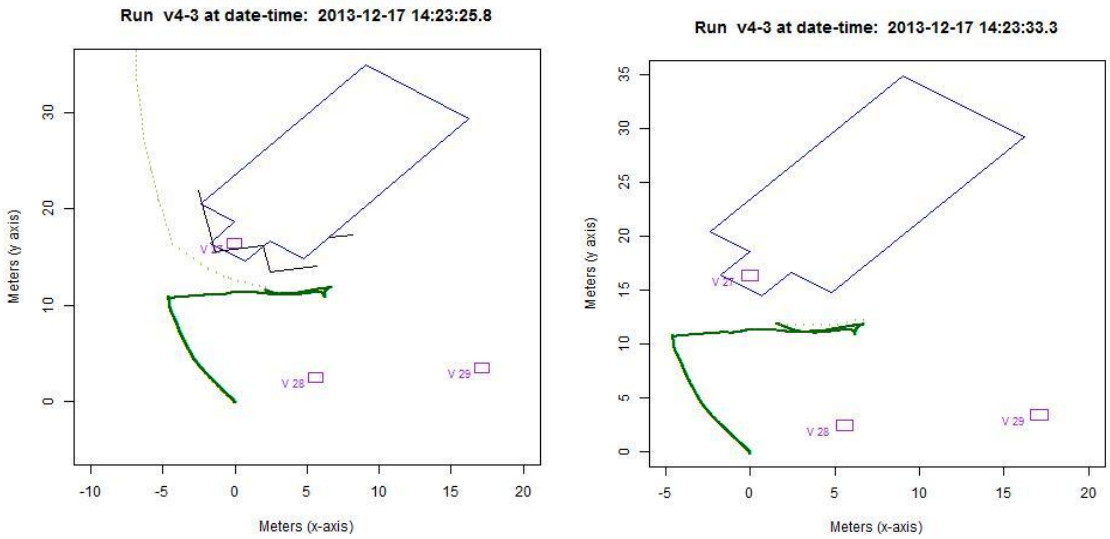

Fig. A-45 More navigation planning from vignette 4, path 3

Analyst comments: The robot violated constraints in a way that is inconsistent with the robot's model of the world.

Table A-44 Assessment of vignette 4, path 3, by researcher

TBS	Navigate left of building to a traffic cone near a car.
Result	Fail
Explanation	[GR] The robot did not see any traffic cones throughout the run and failed to detect any objects near a car. The robot tried to navigate to left of the building but failed to set the correct goal position in xy coordinates.

A.14 Vignette 4, Path 3a, 17 December 2013 at 1436

Table A-45 Human observer scoring for vignette 4, path 3a, during experiment

Aspect	Assessment	Comments
Appropriate navigation	Partial success	Command: "Navigate left of building to a cone near vehicle."
Following constraints	Partial success	Headed in the right direction, then stopped.
Reaching goal	Fail	. . .
Other	. . .	Stopped by operators after "Building flipped".
Overall	F 1	. . .

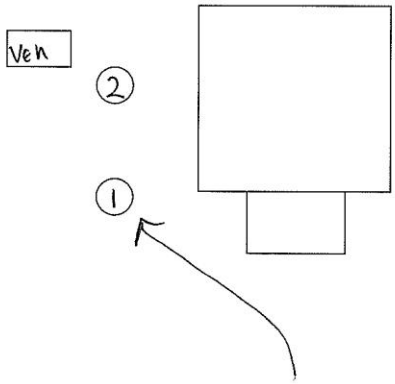

Fig. A-46 Layout for vignette 4, path 3a, 17 December 2013 at 1436

Table A-46 Assessment of vignette 4, path 3a, based on world-model data

Assessed	Time	Code	Comments
Building	14:38:26.7	W, N	Difficult to decide on the W code; see Fig. A-47.
Nav. Plan	14:38:01.3	R	Correct navigation plan, Fig. A48.
Cones	Final	. . .	0 observed, Fig. A-47.
Vehicles	Final	. . .	1 observed but not the one near the cone, Fig. A-47.

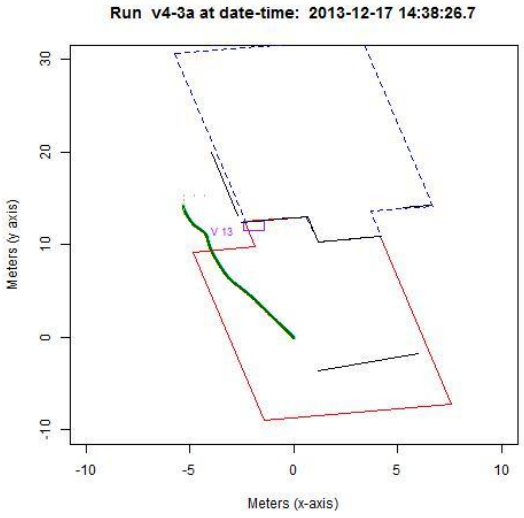

Fig. A-47 World-model reconstruction for vignette 4, path 3a

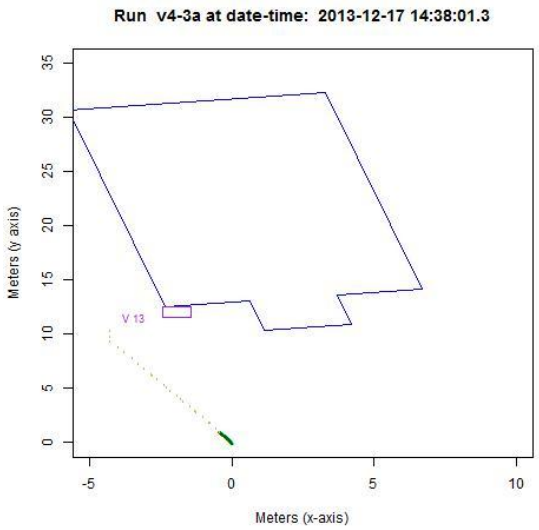

Fig. A-48 World-model reconstruction for vignette 4, path 3a

Table A-47 Assessment of vignette 4, path 3a, by researcher

TBS	Navigate left of building to a traffic cone near a car.
Result	Fail
Explanation	[OD/PR] The robot failed to detect a car; the predicted building was placed on the wrong side, resulting in errors in spatial constraint reasoning.

A.15 Vignette 4, Path 3b, 17 December 2013 at 1455

Table A-48 Human observer scoring for vignette 4, path 3b, during the experiment

Aspect	Assessment	Comments
Appropriate navigation	Partial success	Command: "Navigate left of building to a cone near vehicle."
Following constraints	Fail	Headed left, then turned around and went right.
Reaching goal	Fail	. . .
Other	. . .	Stopped by operator to avoid hitting another robot.
Overall	F 3	. . .

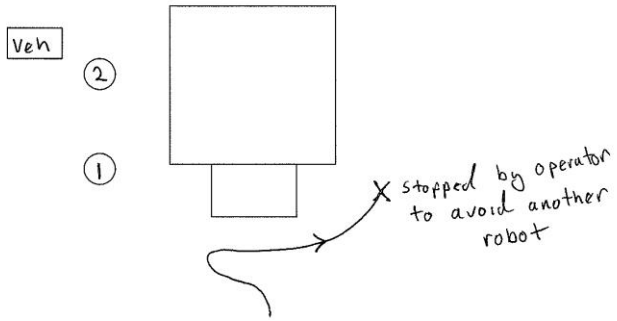

Fig. A-49 Layout for vignette 4, path 3b

Table A-49 Assessment of vignette 4, path 3b, based on world-model data

Assessed	Time	Code	Comments
Building	Final	R	Consistent throughout the run.
Nav. plan	14:56:08.2 14:57:13.8	C	No clear reason for changes between routes, Figs. A-50 and A-51.
Cones	Final	. . .	0 observed, Fig. A-51.
Vehicles	Final	. . .	1 observed but not the one near the cone, Fig. A-51.

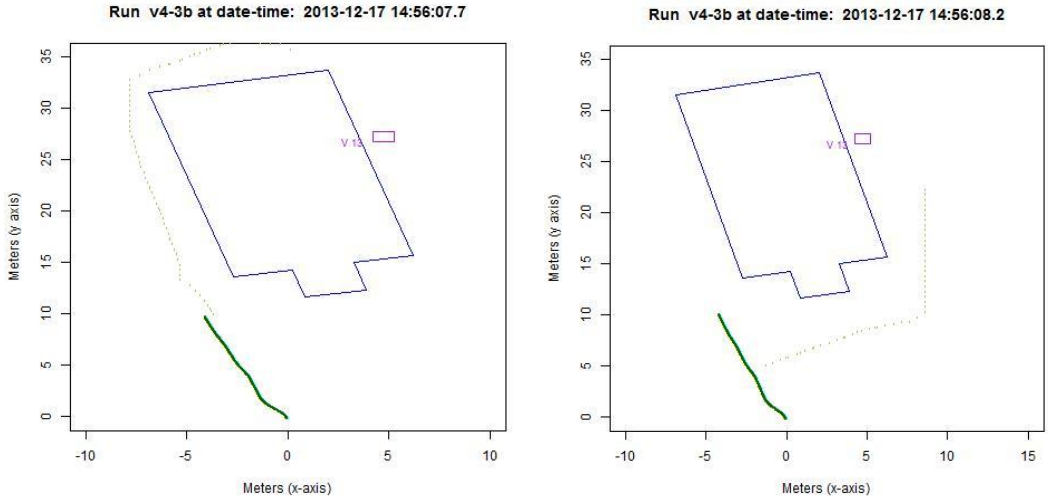

Fig. A-50 World-model reconstruction for vignette 4, path 3b

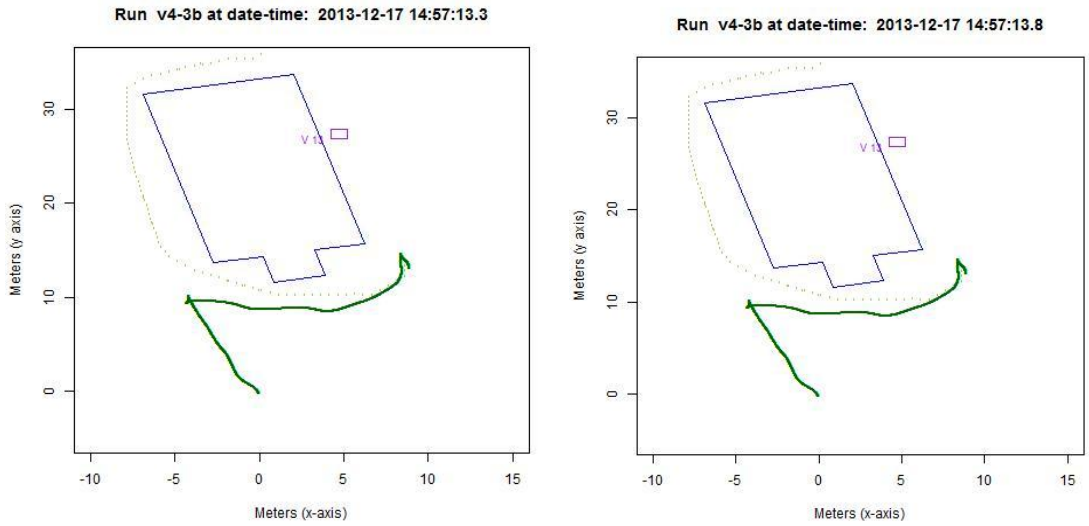

Fig. A-51 World-model reconstruction for vignette 4, path 3b

Analyst comments: The building prediction was consistent throughout the run, but the robot violated navigation constraints.

Table A-50 Assessment of vignette 4, path 3b, by researcher

TBS	Navigate left of building to a traffic cone near a car.
Result	Fail
Explanation	[GR/OD] The robot did not see any traffic cones throughout the run. The robot incorrectly detected a car on the right side of the building. The grounder picked an unknown object in the back of the building first and then chose to go to the other unknown object on the right.

A-15 Vignette 4, Path 3c, 17 December 2013 at 1519

Table A-51 Human observer scoring for vignette 4, path 3c, during the experiment

Aspect	Assessment	Comments
Appropriate navigation	Success	Command: "Navigate left of building to a cone near vehicle."
Following constraints	Success	Went left, far out of the way, and stopped near the rear of the vehicle. Operators claim cone pixels on the back of the vehicle.
Reaching goal	Partial success	Stopped 2-3 m from cone.
Other	. . .	Robot started closer to building than in previous run.
Overall	P 1	. . .

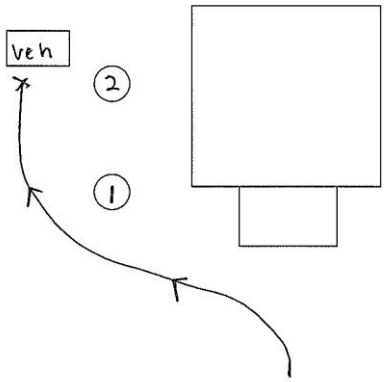

Fig. A-52 Layout for vignette 4, path 3c

Table A-52 Assessment of vignette 4, path 3c, based on world-model data

Assessed	Time	Code	Comments
Building	15:20:35.1	W, N	Fig. A-53
Nav. Plan	. . .	R	Reasonable planning, Fig. A-54.
Cones	Final	. . .	3 observed, Fig. A-55.
Vehicles	Final	. . .	6 observed, Fig. A-55.

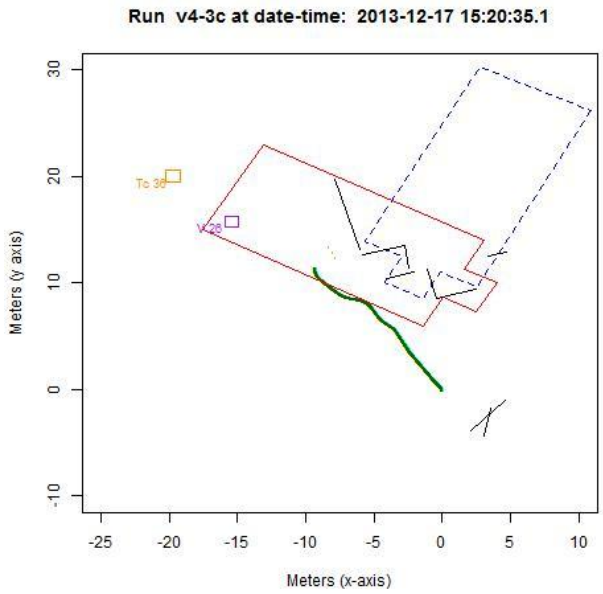

Fig. A-53 Building transformation from vignette 4, path 3c

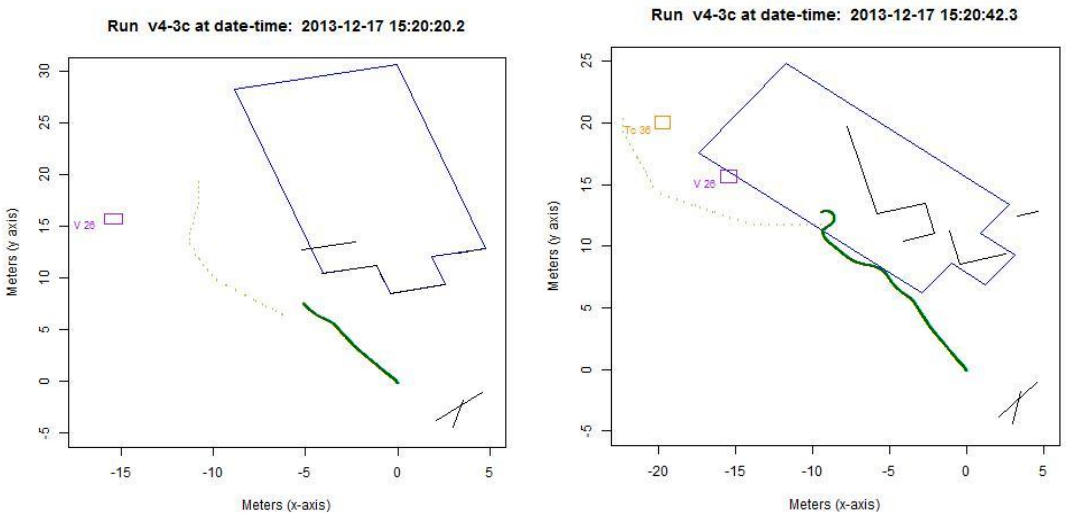

Fig. A-54 Navigation planning from vignette 4, path 3c

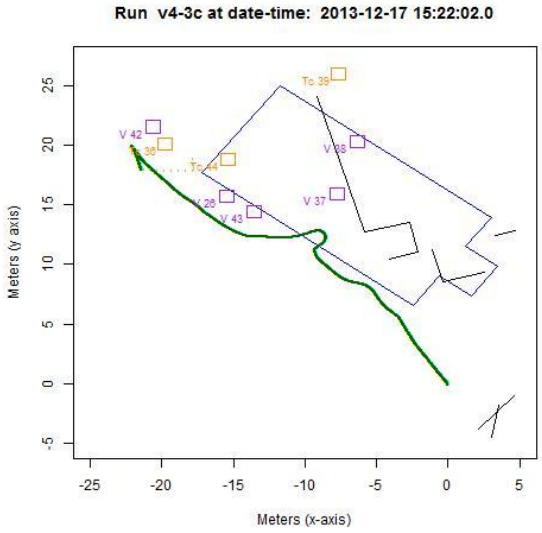

Fig. A-55 Final model from vignette 4, path 3c

Analyst comments: This run was a success, but it is strange that the robot did not seem to see cone 1 until after it had seen cone 2. The intelligence might be able to distinguish between a cone near a vehicle and a cone not near a vehicle, but the multiplicity of false positive vehicles makes this a hard command to execute successfully.

Table A-53 Assessment of vignette 4, path 3c, by researcher

TBS	Navigate left of building to a traffic cone near a car.
Result	Success
Explanation	The robot was able to detect a car on the left side of the building. Initially, the robot aimed at an unknown object on the left of the building. As the robot approached the car, it detected a traffic cone near the car. The robot then re-planned and finished at the correct goal. The wheels were slipping heavily due to snow; resulting drift in micro-LADAR data.

A.15 Vignette 4, Path 3d, 17 December 2013 at 1530

Table A-54 Human observer scoring for vignette 4, path 3d, during the experiment

Aspect	Assessment	Comments
Appropriate navigation	Success	Command: "Navigate left of building to a cone near vehicle."
Following constraints	Success	. . .
Reaching goal	Success	Stopped 1 m from cone.
Other	. . .	Went directly to the cone.
Overall	S	. . .

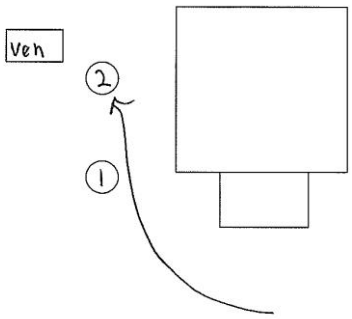

Fig. A-56 Layout for vignette 4, path 3d

Table A-55 Assessment of vignette 4, path 3d, based on world-model data.

Assessed	Time	Code	Comments
Building	15:31:32.0	W	Fig. A-57
Nav. plan	15:31:16.0 15:32:26.0	R	Appropriate planning, Fig. A-58.
Cones	Final	. . .	4 observed, Fig. A-59.
Vehicles	Final	. . .	4 observed (not all shown), Fig. A-59.

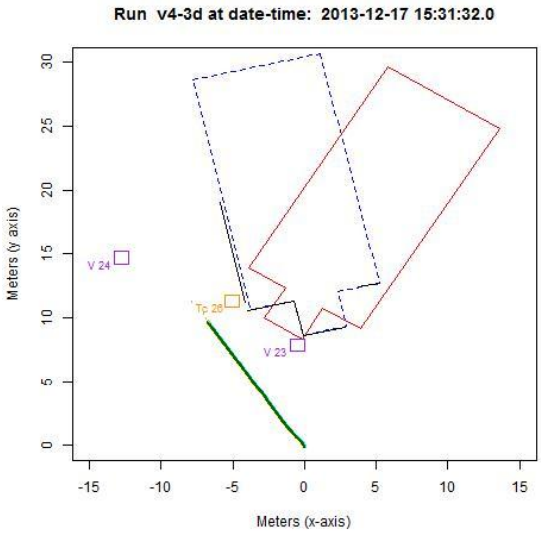

Fig. A-57 Building transformation for vignette 4, path 3d

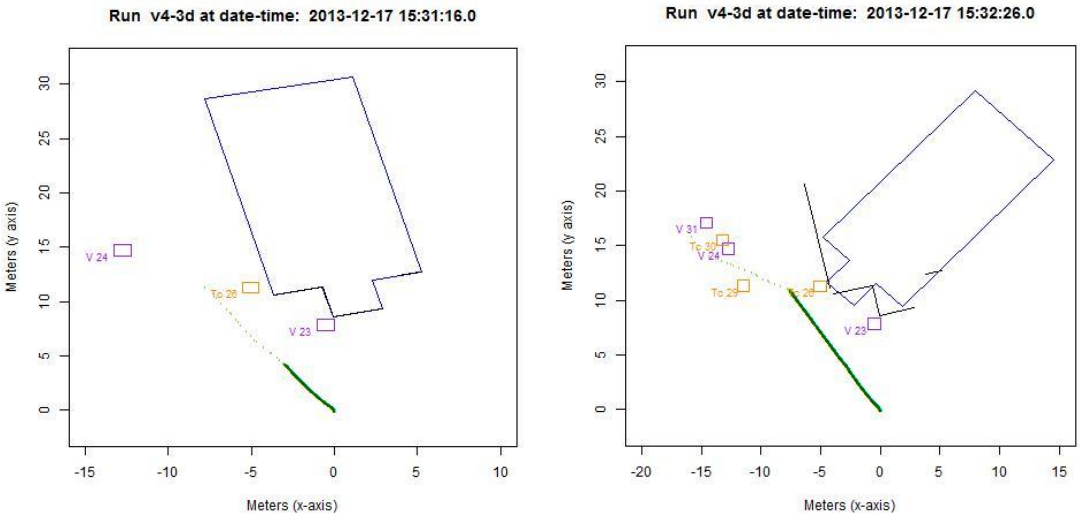

Fig. A-58 Navigation planning from vignette 4, path 3d

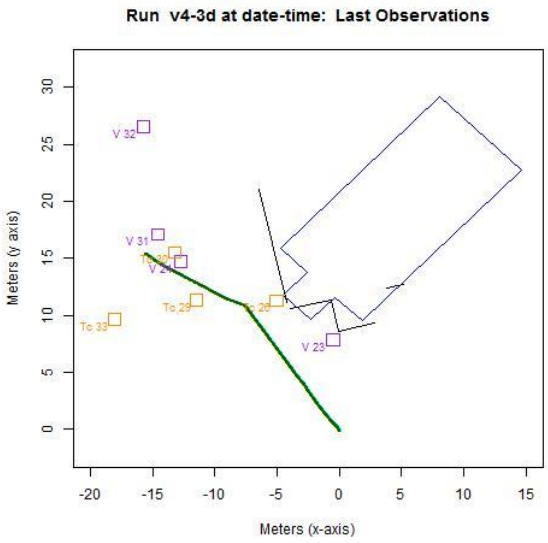

Fig. A-59 Final model for vignette 4, path 3d

Analyst comments: Here it did see a cone, went to it, and then saw a second cone near a vehicle and went to that.

Table A-56 Assessment of vignette 4, path 3d, by researcher

TBS	Navigate left of building to a traffic cone near a car.
Result	Success
Explanation	Consistent behavior with vignette 4, path 3c. [OD] Due to fusion errors, there were several false positive traffic cones.

A.16 Vignette 4, Path 4a, 17 December 2013 at 1550

Table A-57 Human observer scoring for vignette 4, path 4a, during the experiment

Aspect	Assessment	Comments
Appropriate navigation	Fail	Command: "Navigate left of vehicle to a cone near the building."
Following constraints	Fail	Went right of vehicle, down the right side of the building, and then back.
Reaching goal	Fail	. . .
Other	. . .	Stopped on the stairs to avoid an accident.
Overall	F 3	. . .

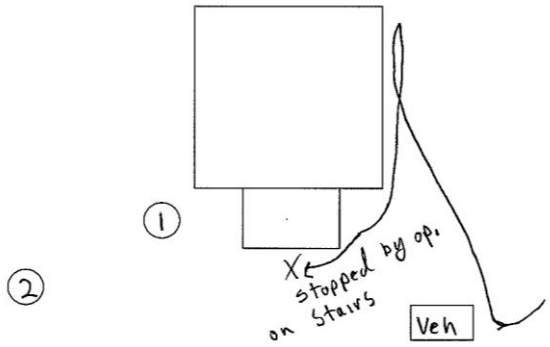

Table A-58 Layout for vignette 4, path 4a

Table A-59 Assessment of vignette 4, path 4a, based on world-model data

Assessed	Time	Code	Comments
Building	15:54:18.9	W	Fig. A-62.
Nav. plan	15:52:41.2	C	Chose a path right of the vehicle, then tried to get back to the front of the building, Figs. A-61 and A-62.
Cones	Final	. . .	1 after run ended, Fig. A-63.
Vehicles	Final	. . .	13 observed but not until after violating the navigation constraint, Fig. A-63.

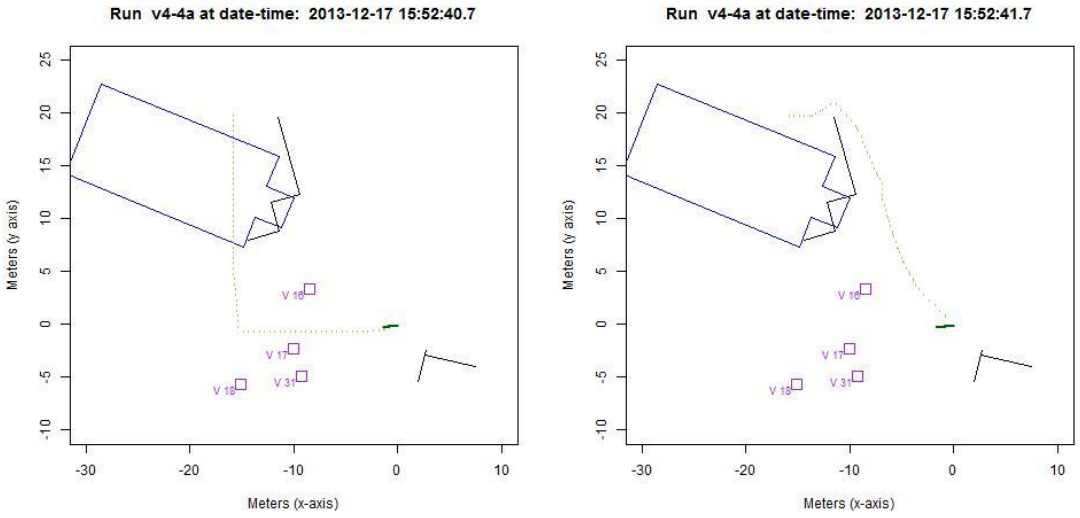

Fig. A-60 Navigation planning from vignette 4, path 4a

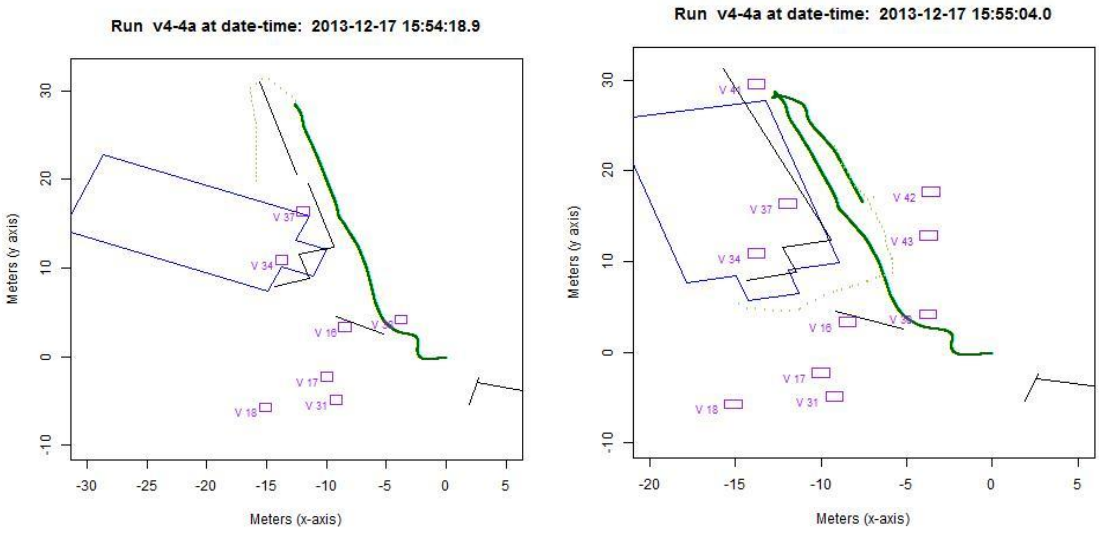

Fig. A-61 Building position from vignette 4, path 4a

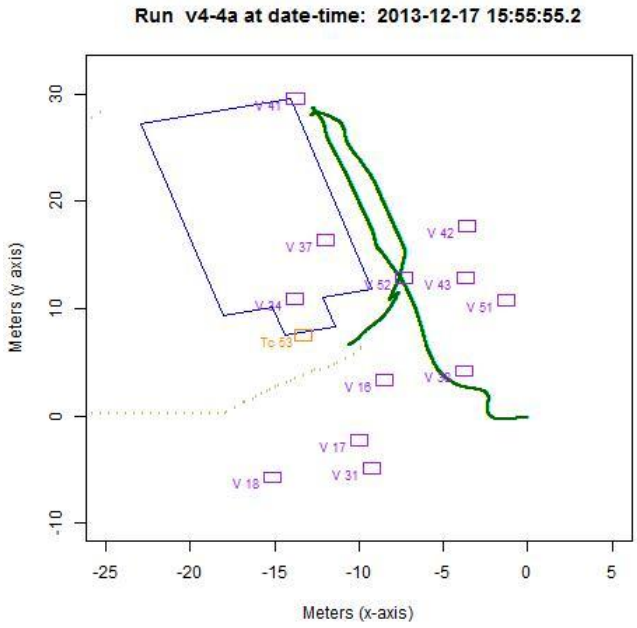

Fig. A-62 Final model from vignette 4, path 4a

Analyst comments: The robot violated the “go left of vehicle” constraint. The robot later planned a path back to the front of the building.

Table A-60 Assessment of vignette 4, path 4a, by researcher

TBS	Navigate left of car to a traffic cone near the building.
Result	Fail
Explanation	[PR/OD] The predicted building had incorrect orientation, and the robot was not able to detect a traffic cone throughout the run. There were also many false positives of cars. The robot attempted to go to an unknown object near the predicted building.

A.17 Vignette 4, Path 4b, 17 December 2013 at 1605

Table A-61 Human observer scoring for vignette 4, path 4b, during the experiment

Aspect	Assessment	Comments
Appropriate navigation	Success	Command: "Navigate left of vehicle to a cone near the building."
Following constraints	Success	Went left of vehicle.
Reaching goal	Success	Stopped within 1 m of cone.
Other	. . .	. . .
Overall	S	. . .

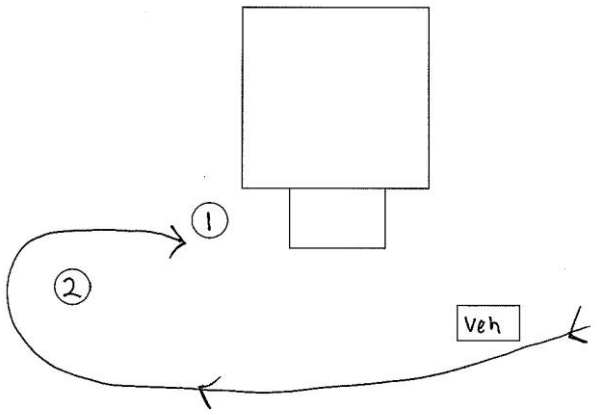

Fig. A-63 Layout for vignette 4, path 4b

Table A-62 Assessment of vignette 4, path 4b, based on world-model data

Assessed	Time	Code	Comments
Building	16:06:39.1	W, N	Fig. A-66.
Nav. plan	16:06:00.8 16:06:45.6	R	Appropriate plans, Fig. A-65.
Cones	Final	. . .	4 observed, Fig. A-67.
Vehicles	Final	. . .	9 observed, Fig. A-67.

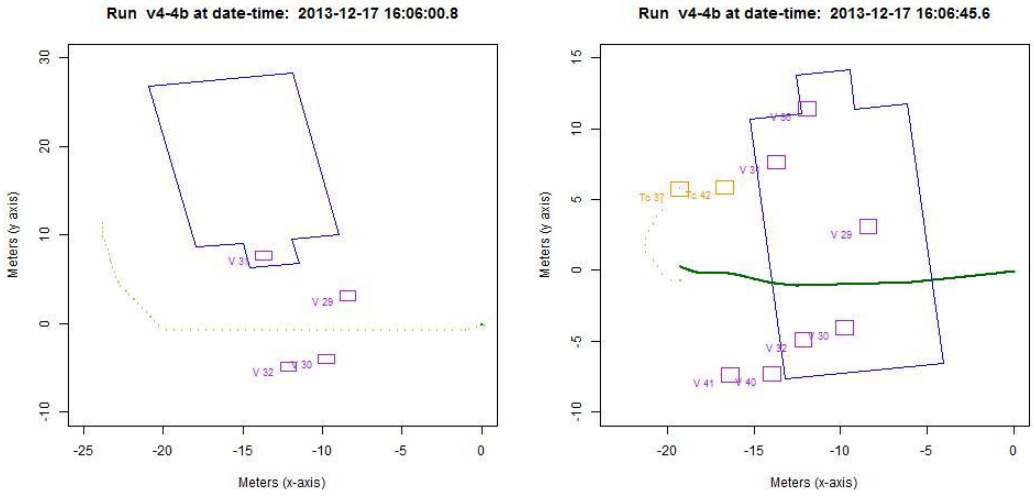

Fig. A-64 Navigation planning from vignette 4, path 4b

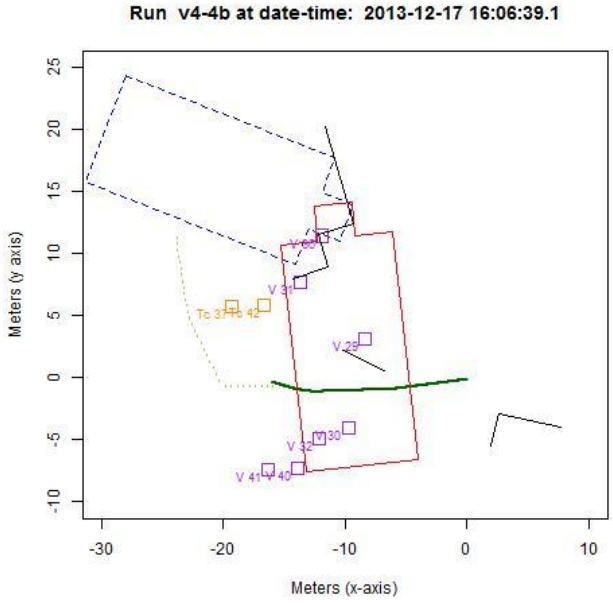

Fig. A-65 Building transformations from vignette 4, path 4b

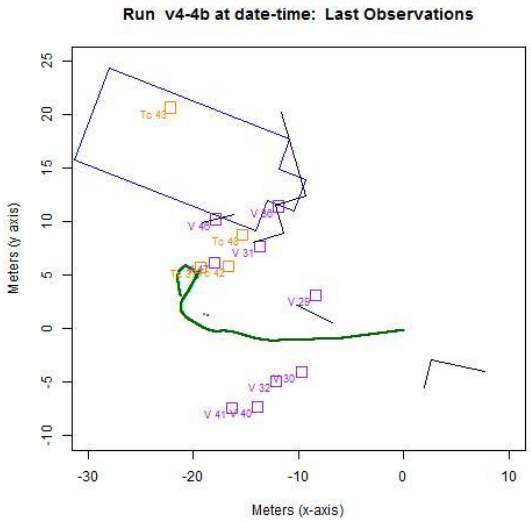

Fig. A-66 Final model from vignette 4, path 4b

Analyst comments: It is not clear why the navigation planner violated the "left of vehicle" constraint in vignette 4, path 4a, but followed it in this one.

Table A-63 Assessment of vignette 4, path 4b, by researcher

TBS	Navigate left of car to a cone near the building.
Result	Success
Explanation	The robot navigated left of a car trying to go to an unknown object near the predicted building. As the robot moved closer to the building, it detected a traffic cone near the building and finished at the correct goal.

1 (PDF) DEFENSE TECHNICAL INFORMATION CTR DTIC OCA

2 (PDF) DIRECTOR US ARMY RESEARCH LAB RDRL CIO LL IMAL HRA MAIL & RECORDS MGMT

1 (PDF) GOVT PRINTG OFC A MALHOTRA

3 (PDF) DIR USARL RDRL CII T B BODT RDRL VTA M CHILDERS C LENNON

INTENTIONALLY LEFT BLANK.