ARL-TR-8352 ● APR 2018

US Army Research Laboratory

Engineering Evaluation and Assessment (EE&A) Report for the Symbolic and Sub-symbolic Robotics Intelligence Control System (SS-RICS)

by Troy Dale Kelley, Eric Avery, and Sean McGhee

NOTICES

Disclaimers

ARL-TR-8352 ● APR 2018

US Army Research Laboratory

Engineering Evaluation and Assessment (EE&A) Report for the Symbolic and Sub-symbolic Robotics Intelligence Control System (SS-RICS)

by Troy Dale Kelley and Eric Avery
Human Research and Engineering Directorate, ARL

Sean McGhee
STG Inc, Reston, VA

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

1. REPORT DATE *(DD-MM-YYYY)*	2. REPORT TYPE	3. DATES COVERED (From - To)
April 2018	Technical Report	1 October 2016–1 March 2018

4. TITLE AND SUBTITLE
Engineering Evaluation and Assessment (EE&A) Report for the Symbolic and Sub-symbolic Robotics Intelligence Control System (SS-RICS)

5a. CONTRACT NUMBER

5b. GRANT NUMBER

5c. PROGRAM ELEMENT NUMBER

6. AUTHOR(S)
Troy Dale Kelley, Eric Avery, and Sean McGhee

5d. PROJECT NUMBER

5e. TASK NUMBER

5f. WORK UNIT NUMBER

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
US Army Research Laboratory
ATTN: RDRL-HRF-D
Aberdeen Proving Ground, MD 21005-5425

8. PERFORMING ORGANIZATION REPORT NUMBER
ARL-TR-8352

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
The US Army Research Laboratory's Human Research and Engineering Directorate has been developing a robotics control architecture called the Symbolic and Sub-symbolic Robotics Intelligence Control System (SS-RICS) since 2006. Since that time, SS-RICS has been the integration platform for many robotics algorithms using a variety of different disciplines from cognitive psychology, computer science, and artificial intelligence. The capabilities of SS-RICS are broad, and an Engineering Evaluation and Assessment was needed to test those operation capabilities in an empirical setting.

15. SUBJECT TERMS
robotics, engineering, assessment, Sub-symbolic Robotics Intelligence Control System, SS-RICS

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
					Troy Dale Kelley
a. REPORT	b. ABSTRACT	c. THIS PAGE	UU	28	19b. TELEPHONE NUMBER (Include area code)
Unclassified	Unclassified	Unclassified			410-278-5869

Standard Form 298 (Rev. 8/98)
Prescribed by ANSI Std. Z39.18

Contents

List of Figures

List of Tables

1. Introduction

The purpose of this report is to provide the results from the Engineering Evaluation and Assessment (EE&A) for the Symbolic and Sub-symbolic Robotics Intelligence Control System (SS-RICS), a robotics control architecture.

The US Army Research Laboratory's Human Research and Engineering Directorate has been developing SS-RICS since 2006 (Kelley 2006). SS-RICS has been the integration platform for many robotics algorithms using a variety of different disciplines from cognitive psychology, computer science, and artificial intelligence. The capabilities of SS-RICS are broad, and an EE&A was needed to test those operation capabilities in an empirical setting.

The development of SS-RICS has been documented in conference papers (Kelley 2006; Avery et al. 2006). The algorithms for episodic memory (Kelley and McGhee 2013; Kelley 2014), learning generalization (Kelley and Long 2010), and improved human–robot interaction (Cassenti et al. 2009, 2012) have been chronicled as well. The system has served as a platform for insights into developmental psychology (Kelley and Cassenti 2011; Kelley 2014) and robot consciousness (Long and Kelley 2009a,b).

SS-RICS is a production-system, goal-oriented, robotics control system based on the cognitive architecture Adaptive Character of Thought–Rational (Anderson and Lebiere 1998). Development of the system has leveraged heavily from biological and human capabilities for navigation, action selection, and memory decay.

2. Mission Need

SS-RICS is currently supporting the Robotics Collaborative Technology Alliance (RCTA) for design and development of algorithms and capabilities for small unmanned ground vehicles (UGVs). The RCTA's current focus is research of small UGVs capable of providing reconnaissance and intelligence in support of a squad.

Several test scenarios have been proposed; the current scenario being researched is to "cover the back of the building", where a UGV is tasked to provide security at the back of a building.

3. System Design

The development of SS-RICS was inspired by the separation of knowledge representation into symbolic and discrete components and sub-symbolic, metric components. SS-RICS was inspired by early work done on cognitive architectures, which were essentially production systems. This led to the development of 3 basic sections for SS-RICS: Working Memory (WM), Perception, and Long-Term Memory (LTM) (Fig. 1).

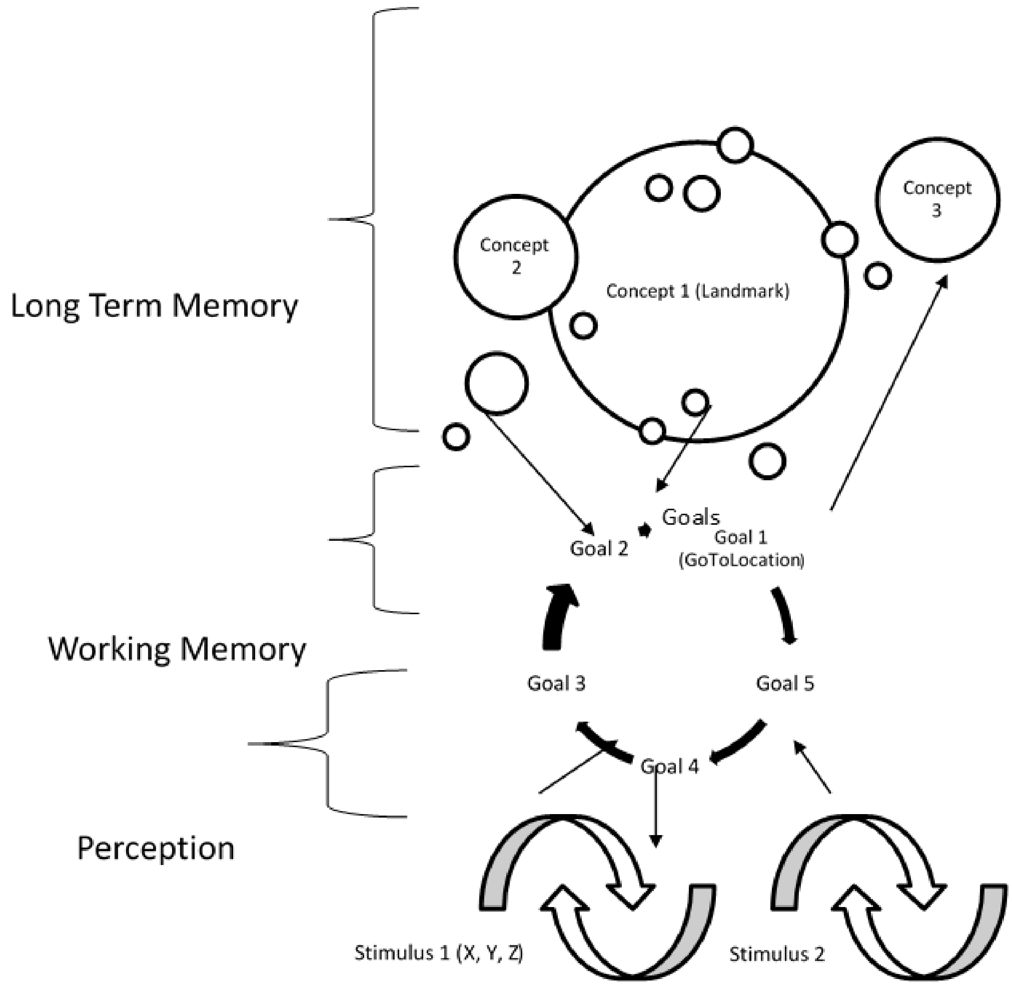

Fig. 1 SS-RICS design. LTM, or Conceptual Memory, contains a symbolic, relational database. WM is a goal stack similar to a production system. Sub-symbolic or metric information is handled by multiple simultaneous algorithms.

At the core of SS-RICS is a production system architecture. The production system simulates the functionality of human WM. The basic unit of knowledge, or WM element, is a fact. Facts are composed of a tuple: name, type, and slot. For example, a fact might contain information such as Location234 color value=red. In this example, Location234 is the name of the fact, the type of the fact is “color”, and the slot, which has the name of “value” is defined as “red”.

Perception is composed of sub-symbolic processors that all run in parallel and feed information to WM. These processors create facts for WM to operate upon depending on the goal within WM. They can perform a variety of tasks, such as search for a color, search for a face, move the robot, and keep track of state information. The important part of the sub-symbolic processor implementation is that the processors all work in parallel to inform WM, and WM selects appropriate information based on goals and subgoals.

LTM is used for conceptual storage, and its implementation is similar to that of a relational database or a semantic web. LTM is built upon ConceptNet from the Massachusetts Institute of Technology (Liu and Singh 2004), but it is anticipated that this will expand as SS-RICS matures. The information in LTM is conceptual, symbolic, and factual information that can be queried to return results to WM. For example, LTM could be queried to determine that a “Soldier is also a Person”. This would allow information that is already known by the system about a person to be applied to Soldiers as well. This is only one example of the kind of query that can be accomplished by LTM.

By integrating Perception with WM and LTM we believe we are capturing the major components of cognition at a functional level. Certainly, many more components of cognition could be represented, and this is not an exhaustive list of the components of cognition within SS-RICS, but it represents the high-level architectural breakout of the system.

SS-RICS is a Windows-based computer program designed to run under the Windows operating system. It was developed using the Microsoft.Net framework. It is capable of running both managed and unmanaged code and is currently being upgraded to run on 64-bit systems and expanded to run across multiple computer processing units.

4. Data Sources

Testing was conducted in multiple phases at Aberdeen Proving Ground in Aberdeen, Maryland. A list of all the algorithms available for testing is contained in Table 1. The first phase of the testing had 4 of the algorithms listed in Table 1: hearing (dictation), face recognition, novelty recognition, and navigation.

Table 1 Experiment breakout for gender, path, noise, and number of runs

Gender	Path	Noise	No. runs
M	A	65 dB	10
F	A	65 dB	10
M	B	80 dB	10
F	B	80 dB	10
M	A	No noise	10
F	B	No noise	10

The first phase had the SS-RICS recognize an operator's face, recognize operator voice commands, navigate to a location without hitting any objects, and return to the home location. The goal was to do this at an 80% success rate, to include completing all of the subtasks. In addition to navigation, the test included higher-level cognitive tasks. During the navigation SS-RICS was required to recognize "interesting" or "novel" locations. The test scenario is described in more detail in Section 6.

5. Accomplishments

- SS-RICS was tested to be capable of responding to voice input using a microphone to execute one step of the production system (i.e., one command). The command set was limited to approximately 20 commands. Commands also included a series of actions that can be triggered by one command; for example, "Go to the Conference Room".
- The voice recognition command set was tested with one woman and one man.
- The voice recognition was tested with a background noise of 65/80 dB of constant white noise.
- SS-RICS was tested to be capable of recognizing a specific operator's face following a period of learning and training. This learning and training period took no more than 5 min.
- SS-RICS was capable of navigating to a specific location, without hitting any objects, using a previously defined map.
- SS-RICS was capable of returning to a "home" location after each completed navigation procedure without hitting any objects.
- SS-RICS was capable of recognizing interesting or novel points in the navigation environment.

The remainder of this report focuses on details of the operational testing and comparisons of facial recognition algorithms.

6. Operational Testing

Testing was a within-subject 2 (gender) × 2 (paths) × 3 (noise level) design. Gender testing was conducted because previous pilot tests showed differences in voice recognition based on gender. Path testing was conducted to test the navigation algorithms (simultaneous localization and mapping [SLAM]) and the landmark designation algorithms. Background noise was added to see what effect this had on voice recognition. Each noise level was run 10 times per gender, yielding 60 total runs.

Two paths were chosen for testing (Paths A and B) of relatively equal length and complexity. At the end of each run the completions were noted: complete or not complete. The path was considered complete when the robot reached the final destination point and then returned to the home location.

Two noise levels were tested: 65 and 80 dB of white noise, with a baseline of no noise. Decibels levels were measured prior to the path runs (A, B). Noise levels were reversed for each gender. Males started with no noise and increased. Females started with 80 dB and decreased. Voice recognition was true or false for whether or not the command was recognized on the first utterance.

Face recognition was tested as either true or false at the beginning of each run, with the operator showing their face following some period when their face was not visible.

7. Face Recognition Algorithm Comparison

Three models of face recognition algorithms are available in SS-RICS: EigenFaces, FisherFaces, and Local Binary Pattern Histograms (LBPHs). These were all provided by the package openCV (www.openCV.org). We ran initial tests on all 3 face recognition algorithms on a variety of subjects, male and female, and then settled on LBPHs. The algorithm comparison follows.

Nine subjects (6 men, 3 women) were used to train each model with 10 still video images of each subject comprising the inner mask of the face from approximately the eyebrows to the lower lip and largely without ears or extraneous background or clothing. Each image was approximately 130 × 130 resolution and scaled to 200 × 200 for training the model.

Each face used for classification (all from the pool of 9) was captured by the same software mechanism that captured the faces for training at a resolution of approximately the same 130 × 130 and scaled up to 200 × 200 for both training and classification.

All faces for training and classification had the same lighting parameters and were captured in the same locale. All images were grayscaled prior to training or classification (each model requires this).

Each classification test involved capturing the candidate's face for a total of 100 frames and using the 3 models to classify it. The results were compared with the ground truth of who the candidate actually was.

Of the 9 training subjects, subjects 2, 5, 7, 8, and 9 were available for further testing and training (Tables 2–6).

Table 2 Testing subject 2

Testing subject 2	EigenFaces	FisherFaces	LBPHs
Subject 1	. . .	. . .	. . .
Subject 2	28	73	76
Subject 3	0	11	2
Subject 4	0	0	0
Subject 5	41	0	0
Subject 6	31	0	22
Subject 7	0	0	0
Subject 8	0	16	0
Subject 9	0	0	0

Percentage of correct classifications by model:
LBPHs = 76%
FisherFaces = 73%
EigenFaces = 28%

Table 3 Testing subject 5

Testing subject 5	EigenFaces	FisherFaces	LBPHs
Subject 1	0	0	0
Subject 2	0	0	0
Subject 3	0	15	0
Subject 4	0	0	0
Subject 5	99	85	100
Subject 6	1	0	0
Subject 7	0	0	0
Subject 8	0	0	0
Subject 9	0	0	0

Percentage of correct classifications by model:
LBPHs = 100%
FisherFaces = 85%
EigenFaces = 99%

Table 4 Testing subject 7

Testing subject 7	EigenFaces	FisherFaces	LBPHs
Subject 1	0	0	0
Subject 2	0	0	0
Subject 3	0	0	3
Subject 4	0	2	1
Subject 5	0	0	0
Subject 6	0	0	18
Subject 7	100	96	78
Subject 8	0	0	0
Subject 9	0	0	0

Percentage of correct classifications by model:
LBPHs = 78%
FisherFaces = 96%
EigenFaces = 100%

Table 5 Testing subject 8

Testing subject 8	EigenFaces	FisherFaces	LBPHs
Subject 1	2	0	0
Subject 2	0	0	0
Subject 3	0	0	0
Subject 4	0	0	0
Subject 5	0	0	1
Subject 6	0	0	0
Subject 7	0	0	0
Subject 8	98	100	99
Subject 9	0	0	0

Percentage of correct classifications by model:
LBPHs = 99%
FisherFaces = 100%
EigenFaces = 98%

Table 6 Testing subject 9

Testing subject 9	EigenFaces	FisherFaces	LBPHs
Subject 1	0	0	0
Subject 2	0	0	1
Subject 3	6	24	1
Subject 4	0	4	0
Subject 5	45	0	4
Subject 6	1	0	4
Subject 7	0	0	0
Subject 8	12	0	0
Subject 9	36	73	90

Percentage of correct classifications by model:
LBPHs = 90%
FisherFaces = 73%
EigenFaces = 36%

The model that worked best overall (had the majority of correct classifications over the test period) was the LBPHs. FisherFaces scored second most accurate and EigenFaces third. So we chose the LBPHs method to use for these tests' facial recognition requirement.

8. Test Results

8.1 Overall Reliability

Overall, for the 4 algorithms tested (navigation, landmark designation, voice recognition, and face recognition) the reliability rate was very high for both genders. Overall reliability was above 96% for 60 runs (Fig. 2).

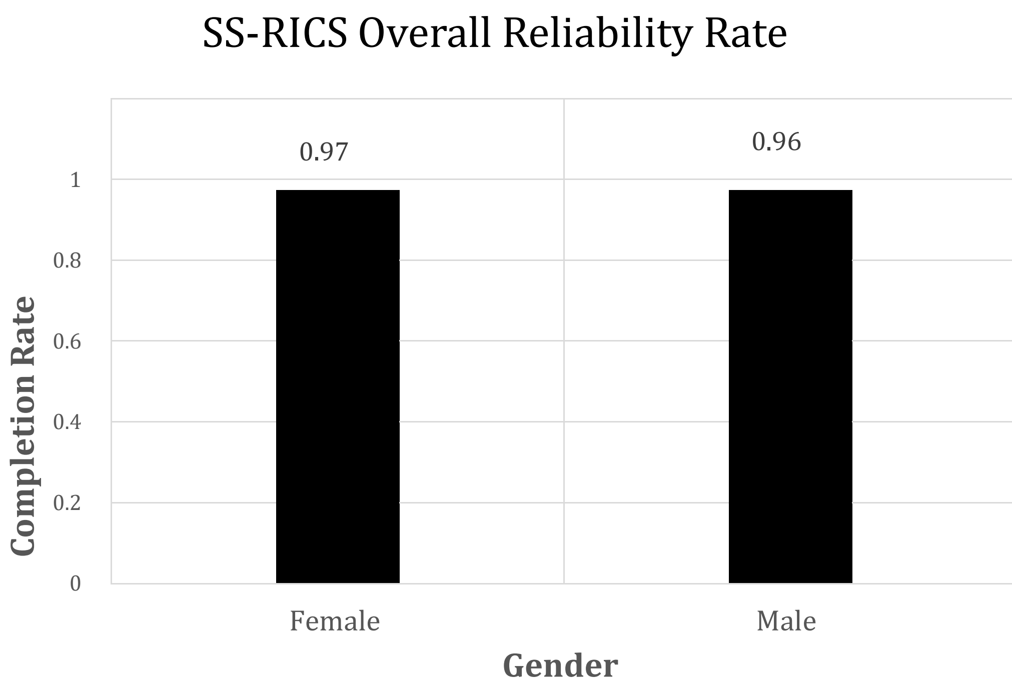

Fig. 2 Overall SS-RICS reliability rate for both genders

8.2 Navigation

For the test procedure, SS-RICS was commanded by the operator to proceed to either of 2 goals: a mannequin or a street sign. The path to each location was approximately equal in distance (~8 m). The goals were alternating over 60 runs, with 30 runs using a male operator and 30 runs using a female (see Table 1).

SS-RICS actually performed a navigation task more than 60 times since the operator had to return the robot to the home location at the beginning of each run. In other words, a typical run would include going to the either the mannequin or the street sign then returning to the home location.

Results from each navigation task were coded as a binary event: the robot made it to the end location or it did not. These were coded at the end of each run.

Overall, there was a 96% success rate for both men and women for SS-RICS over the course of 60 runs (Fig. 3).

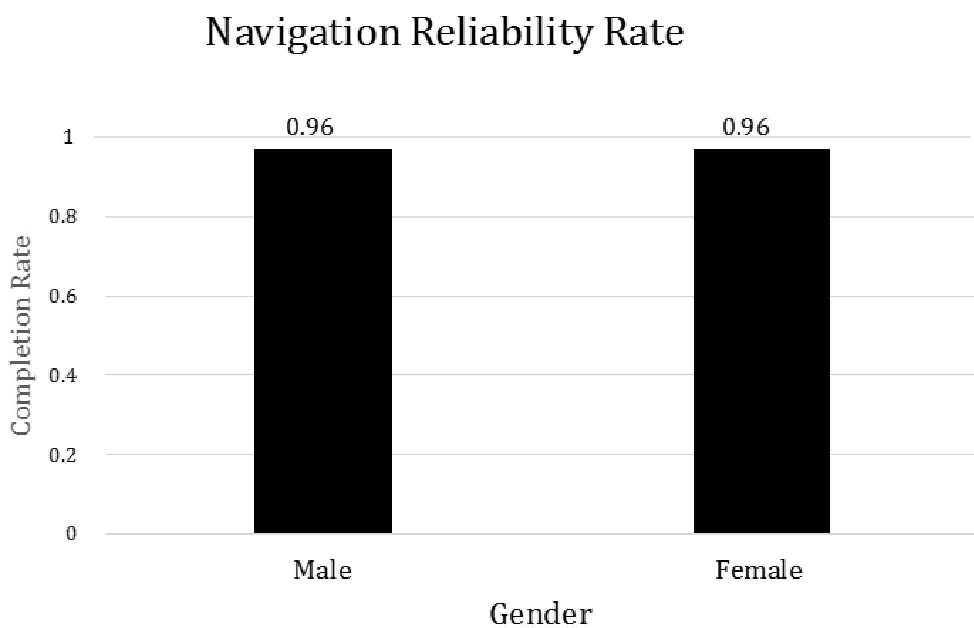

Fig. 3 Navigation reliability rate

8.3 Voice Recognition

For the voice recognition test procedure, SS-RICS was commanded by the operator to navigate to a 1 of 2 goals. This command was given by the operator as a voice input using a wireless microphone mounted on the operator's face in front of the operator's mouth. Before the start of the trials, the operator's voice (both male and female) was trained using the Microsoft's voice recognition system, which is integrated into SS-RICS to allow for voice commands to be given to the robot (Fig. 4). The voice training period took approximately 10 min each to complete for both the male and female operators.

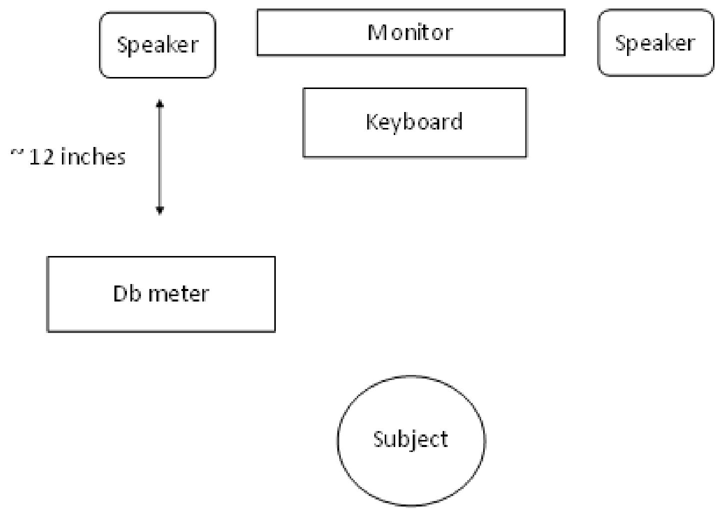

Fig. 4 Overhead view of decibel meter placement

Each operator was tested under 3 levels of background noise: no noise, 65 dB, and 80 dB (see Table 1). The background noise was generated by SS-RICS and sent through the computer speakers. A decibel meter was used to measure the average decibel output from the computer speakers. The decibel meter was placed approximately 12 inches from the computer speakers producing the background noise. The overall recognition rate was 96% for males and 100% for females (Fig. 5).

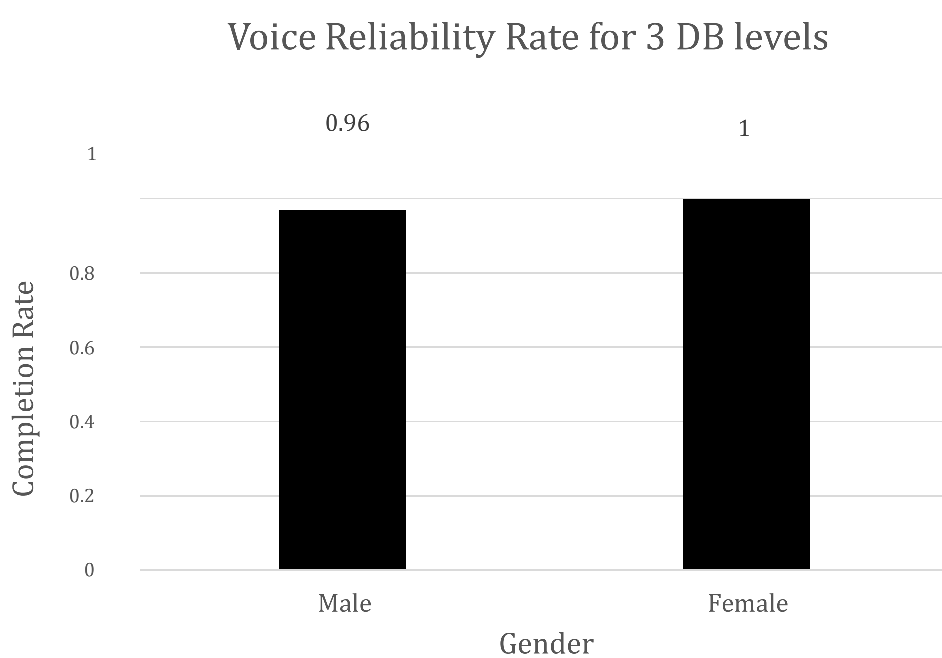

Fig. 5 Voice reliability rate for 3-dB levels

8.4 Face Recognition

Face recognition was conducted before each run using the LBPHs method. Each operator was trained by the SS-RICS LBPHs face recognition algorithm. Training took less than 10 min to complete for each subject.

At the beginning of each run, the operator would stand out of the view of robot's camera so that SS-RICS could not see the operator's face. Once the run began, the operator would move in front of the robot's camera so that SS-RICS could identify the operator.

Since there is an oscillation in the face recognition algorithm over each frame, a set of 100 frames was collected to determine the identity of the operator over a sequence of frames. These 100 frames were then averaged to produce a final determination of the identity of the operator. SS-RICS had a collection of 11 possible operators with which to match to the current operator.

Results were remarkably high and higher than the test and evaluation goal of 80%. The average was 99% each for both the female and male operator (Table 7).

Table 7 Results for facial recognition

Test subject "Bill"	
100%	29 runs
98%	1 run
Average: 99.93%	
Test subject "Kristin"	
100%	23 runs
99%	1 run
98%	2 runs
96%	2 runs
93%	1 run
92%	1 run
Average: 99%	

8.5 Landmark Designation

The landmark designation algorithm finds locations on a map that can be used for landmarks as reference points if the metric map changes. One disadvantage of SLAM and other metric mapping techniques is that each point on a metric map is just as important as every other point, so there are no landmarks that can be used for navigation in a dynamic environment. If the small metric details of a map change, the landmarks should still be present.

Our definition of a landmark is an important location, relative to all other locations, in the map that facilitates future navigation and localization. For example, in our test scenarios, the home location and each one of the goal locations would be considered landmarks (Fig. 6).

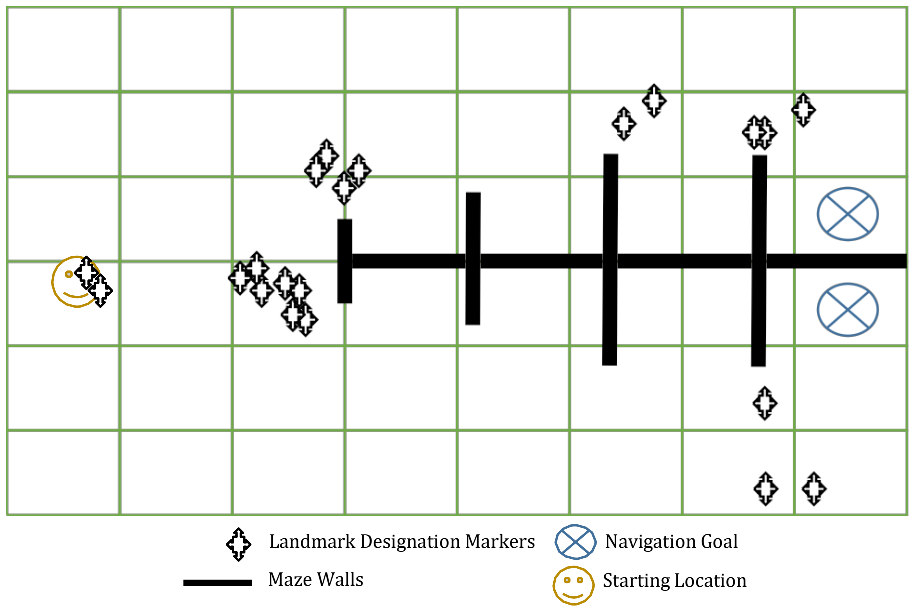

Fig. 6 Landmark designations

Testing was used to determine if SS-RICS would designate the home location and each one of the goal locations (the mannequin and the sign) as landmarks. This landmark designation algorithm has been detailed in a previous report (Kelley and McGhee 2013).

The landmark designation algorithm is stochastic. Over the course of the same path, SS-RICS might designate a landmark at slightly different locations. While the algorithm is stochastic, it responds within a normal distribution of values. Thus, the algorithm was tested to see whether or not it was relatively stable across the testing area for landmark designation. As previously mentioned, the algorithm should designate the home location and the 2 goal locations as landmarks. As can be seen in Fig. 6, SS-RICS designated the home location, the goal locations, and other critical navigation areas (the protrusions from the maze) as landmarks over 60 runs, which is an acceptable level of performance.

Since the evaluation of this version of the algorithm, the algorithm has been improved, so that when the robot travels over a previously designated area, it retrieves existing landmarks. This strengthens their activations (Anderson and Lebiere 1998), thereby reducing the stochasticity of the algorithm and reducing the number of landmark sites.

9. Conclusions

SS-RICS performed well over the course of 60 runs, testing 4 major algorithms, using 3 levels of white background noise, 2 path designations, and 2 gender types as operators. As can be seen in Fig. 7 the reliability percentages were higher than the previously set goal of 80%. Other algorithms within SS-RICS will be tested in the future to ascertain there overall level of reliability.

Testing with multiple operators, of different genders, ethnicities, and socio-economic backgrounds might cause the system more errors and will be the focus of future evaluations.

See the Appendix for future algorithm test scenarios.

10. References

Anderson JR, Lebiere C. The atomic components of thought. Mahwah (NJ): Lawrence Erlbaum Associates Publishers; 1998.

Avery E, Kelley TD, Davani D. Using cognitive architectures to improve robot control: integrating production systems, semantic networks, and sub-symbolic processing. Proceedings of the 15th Annual Conference on Behavioral Representation in Modeling and Simulation (BRIMS); 2006 May 15–19; Baltimore, MD.

Cassenti DN, Kelley TD, Swoboda JC, Patton DJ. The effects of communication style on robot navigation performance. Proceedings of the 53rd Human Factors and Ergonomics Society Annual Meeting; 2009; San Antonio, TX. p. 359–363.

Cassenti DN, Kelley TD, Yagoda RE, Avery E. Improvements in robot navigation through operator speech preferences. Paladyn. 2012;3(2):102–111.

Kelley TD. Developing a psychologically inspired cognitive architecture for robotic control: the symbolic and sub-symbolic robotic intelligence control system (SS-RICS). International Journal of Advanced Robotic Systems. 2006;3(3):219–222.

Kelley TD. Robotic dreams: a computational justification for the post-hoc processing of episodic memories. International Journal of Machine Consciousness. 2014;6(2):109–123.

Kelley TD, Cassenti D. Theoretical explorations of cognitive robotics using developmental psychology. New Ideas in Psychology. 2011;29(3):228–234.

Kelley TD, Long LN. Deep Blue cannot play checkers: the need for generalized intelligence for mobile robots. Journal of Robotics. 2010:1–8.

Kelley TD, McGhee S. Combining metric episodes with semantic event concepts within the symbolic and sub-symbolic robotics intelligence control system. SPIE 8756, Multisensor, Multisource Information Fusion: Architectures, Algorithms, and Applications; 2013 May 29. p. 87560.

Liu H, Singh P. ConceptNet – a practical commonsense reasoning tool-kit. BT Technology Journal. 2004;22(4):211–226.

Long LN, Kelley TD. A review of consciousness and the possibility of conscious robots. Journal of Aerospace Computing, Information and Communication. 2009a;6(6):69–91.

Long LN, Kelley TD. The requirements and possibilities of creating conscious systems. Presented at the AIAA InfoTech@Aerospace Conference; 2009b; Seattle, WA. AIAA Paper No.: 2009–1949.

Appendix. List of Algorithms for Possible Testing

Table A-1 presents algorithms for future testing and includes references to the Adaptive Resonance Theory (ART) of Neural Networks.[1]

Table A-1 List of algorithms for possible testing

Algorithm	Description	Batch or iterative	Testing
Face instance	Recognize a specific face given training.	Iterative	YES
Face recognition	Recognize any person's face.	Batch	YES
Face (left eye)	Recognize any person's left eye.	Batch	NO
Face (right eye)	Recognize any person's right eye.	Batch	NO
Face (profile)	Recognize any person's profile.	Batch	NO
Face (mouth)	Recognize any person's mouth.	Batch	NO
Body (full)	Recognize any person's body.	Batch	NO
Body (upper)	Recognize any person's upper body.	Batch	NO
Body (lower)	Recognize any person's lower body.	Batch	NO
Motion (turning)	Robot will turn a specific degree.	Iterative	YES
Motion (forward/backward)	Robot will move forward or backward a certain number of meters.	Iterative	YES
Motion (obstacle avoidance)	Robot will avoid objects in the front.	Iterative	YES
Speech	Robot will speak written text.	Iterative	YES
Hearing (free)	Robot will listen to and convert a spoken sentence into text. Any words can be used.	Iterative	NO
Hearing (dictation)	Robot will listen to certain key words within a command set.	Iterative	YES
Natural language	Robot will parse a sentence for parts of speech (noun, verb, adverb).	Iterative	NO

[1] Carpenter GA, Grossberg S. Adaptive resonance theory. In: Arbib MA, editor. The handbook of brain theory and neural networks, 2nd ed. Cambridge (MA): MIT Press; 2003. p. 87–90.

Table A-1 List of algorithms for possible testing (continued)

Algorithm	Description	Batch or iterative	Testing
Navigation	Robot will use a pre-existing map to navigate to certain predefined locations.	Iterative	YES
Navigation (Simultaneous Localization and Mapping)	Robot can build a map and simultaneously localize within the map.	Iterative	YES
Choice	Robot can use a stochastic methodology for making choices in an environment.	Iterative	NO
Vision (template match)	Can identify a snippet within a larger visual field.	Batch	NO
Vision (image correlation)	Can correlate the current scene with a saved scene and return the correlation value	Batch	NO
Vision (color)	Can identify colors	Batch	NO
Vision (movement)	Can identify that an object is moving	Batch	NO
Vision (tracking learning detection)	Can track objects identified with a bounding box	Batch	NO
Vision (salience)	Can identify the most salient item in a scene (based on psychological theory)	Batch	NO
Vision (salience)	Can segment an object from a background	Batch	NO
Concepts	Can relate concepts to each other (how similar is a dog and a cat?)	Iterative	NO
Novelty recognition	Can identify new streams of information as being novel from previous streams of information	Batch	YES
Episodic learning	Can learn from experience using previous episodes so that the robot can anticipate events	Iterative	NO
Neural networks	Can recognize laser data using neural networks	Batch	NO

Table A-1 List of algorithms for possible testing (continued)

Algorithm	Description	Batch or iterative	Testing
Neural networks (ART)	Can recognize laser data using ART neural networks	Batch	NO
Memory decay (standard)	The robot has memory decay algorithms similar to human memory decay algorithms	Iterative	NO
Memory decay (habit)	The robot has memory decay algorithms similar to low level neurons which exhibit habituation	Iterative	NO
Production system	The robot has a production system used for selecting behaviors	Iterative	YES

List of Symbols, Abbreviations, and Acronyms

ART	Adaptive Resonance Theory
EE&A	Engineering Evaluation and Assessment
LBPHs	Local Binary Pattern Histograms
LTM	Long-Term Memory
RCTA	Robotics Collaborative Technology Alliance
SLAM	Simultaneous Localization and Mapping
SS-RICS	Symbolic and Sub-symbolic Robotics Intelligence Control System
UGV	unmanned ground vehicle
WM	Working Memory

1 DEFENSE TECHNICAL
(PDF) INFORMATION CTR
DTIC OCA

2 DIR ARL
(PDF) IMAL HRA
RECORDS MGMT
RDRL DCL
TECH LIB

1 GOVT PRINTG OFC
(PDF) A MALHOTRA

1 ARL
(PDF) RDRL HRB B
T DAVIS
BLDG 5400 RM C242
REDSTONE ARSENAL AL
35898-7290

8 ARL
(PDF) SFC PAUL RAY SMITH CENTER
RDRL HRO COL H BUHL
RDRL HRF J CHEN
RDRL HRA I MARTINEZ
RDRL HRR R SOTTILARE
RDRL HRA C A RODRIGUEZ
RDRL HRA B G GOODWIN
RDRL HRA A C METEVIER
RDRL HRA D B PETTIT
12423 RESEARCH PARKWAY
ORLANDO FL 32826

1 USA ARMY G1
(PDF) DAPE HSI B KNAPP
300 ARMY PENTAGON
RM 2C489
WASHINGTON DC 20310-0300

1 USAF 711 HPW
(PDF) 711 HPW/RH K GEISS
2698 G ST BLDG 190
WRIGHT PATTERSON AFB OH
45433-7604

1 USN ONR
(PDF) ONR CODE 341 J TANGNEY
875 N RANDOLPH STREET
BLDG 87
ARLINGTON VA 22203-1986

1 USA NSRDEC
(PDF) RDNS D D TAMILIO
10 GENERAL GREENE AVE
NATICK MA 01760-2642

1 OSD OUSD ATL
(PDF) HPT&B B PETRO
4800 MARK CENTER DRIVE
SUITE 17E08
ALEXANDRIA VA 22350

ABERDEEN PROVING GROUND

12 ARL
(PDF) RDRL HR
J LOCKETT
P FRANASZCZUK
K MCDOWELL
K OIE
RDRL HRB
D HEADLEY
RDRL HRB C
J GRYNOVICKI
RDRL HRB D
C PAULILLO
RDRL HRF A
A DECOSTANZA
RDRL HRF B
A EVANS
RDRL HRF C
J GASTON
RDRL HRF D
A MARATHE
T KELLEY