Improving robot situational awareness through commonsense: Side-stepping incompleteness and unsoundness

Sanford T. Freedman and Julie A. Adams

Abstract— **Robots are continuously being deployed in ever challenging domains with tasks of increasing complexity. With these mounting challenges, robots will need to gain a better understanding of their environment and the status of surrounding dynamic obstacles. Human situational awareness may provide robots with this greater understanding. Presently, no straightforward approach is known for achieving human level situational awareness within robotic systems.**

This paper presents an initial exploration into the possibilities of utilizing commonsense reasoning for the explicit purpose of fostering robot situational awareness. The paper opens with a description of both situational awareness and commonsense reasoning and some of the challenges involved in incorporating these abilities into robotic systems. A possible system is described for improving robot situational awareness through the use of commonsense reasoning. The paper concludes with an initial proof of concept evaluation of the proposed system.

I. INTRODUCTION

Future robotic systems will be asked to perform complex real-time tasks currently performed by humans. Success at these endeavors will require robots to possess a comprehensive collection of skills and knowledge, spanning vision processing, logical reasoning, analogical reasoning, and social conventions. While robots may not require an identical skill set and list of capabilities, human capabilities can provide important inspiration and suggestions into the research of improving robotic capabilities.

Even within domains currently confronting robots, a tremendous volume of information is available, so much so, that processing it in its entirety can be impossible. Further compounding the problem, information required for successful operation may be obscured and unavailable. Similar issues confront humans attempting to develop and maintain Situational Awareness (SA). A phenomena with numerous definitions [1], SA can loosely be considered an understanding of the environment and effects of possible future actions. Humans are able to utilize a wide array of techniques and skills in order to maintain SA [2], which include mental models, schemas, directed attention, automaticity, bottom-up processing, and top-down processing. Humans aiming to maintain SA can also utilize commonsense reasoning.

Commonsense reasoning describes humans' ability to handle basic tasks and understand concepts often referred to as obvious or intuitive [3]. Commonsense provides defaults and generally plausible inferences that aid humans in understanding the world and their actions.

This paper proposes that commonsense reasoning can play a vital role in realizing robotic SA. Defaults and inferences generated by commonsense reasoning can provide data necessary to overcome the inherent unobservability of many real world robotic tasks and domains. Additionally, defaults can aid robots in directing attention to the most probable solutions, reduce computational and sensor load, and allow more time to be spent generating a comprehensive understanding of the current situation. An initial exploration into the effects and benefits to robotic SA by incorporating commonsense reasoning is presented.

II. BACKGROUND

SA and commonsense have received considerable attention across many diverse disciplines. This section provides a brief review of previous efforts.

A. Situational Awareness

SA has received considerable attention in the human factors community [1], [4], [5]. Endsley defines SA as "the perception of the elements in the environment within a volume of time and space, the comprehension of their meaning, and the projection of their status in the near future" [6].

While often misinterpreted as a process (or a set of processes), SA is a grouping and a description of processes working to form human cognition. The portions of these processes that handle incoming data and then compute higher level constructs can be grouped together by SA. Endsley partitioned SA into three levels [7]. Level 1 involves acquiring raw information from the environment. Level 2 consists of merging data from level 1 into an understanding of the current environment situation. The final level utilizes levels 1 and constructs from level 2 to predict the future states of the environment and possible actions to undertake.

B. Commonsense Reasoning

Operating within supposedly simple domains and completing apparently trivial tasks requires a tremendous quantity of knowledge and reasoning capabilities. Children playing with blocks may require the ability to reason over physical, social, mental, bodily, visual, spatial, and tactile domains [8]. Even placing a simple telephone call may require an equally long list of dependent abilities [3]. Despite these daunting lists of potentially required skills and knowledge, humans generally complete and excel at these tasks. Much of this ability can be

S. T. Freedman is with the Department of Electrical Engineering and Computer Science, Vanderbilt University, Nashville, TN 37235, USA sandytf@alumni.usc.edu

J. A. Adams is with the Department of Electrical Engineering and Computer Science, Vanderbilt University, Nashville, TN 37235, USA julie.a.adams@vanderbilt.edu

attributed to commonsense, "the things that we expect other people to know and regard as obvious" [3].

Robotic systems frequently utilize collections of subsystems, each dedicated to a single form of reasoning or subset of tasks. These subsystems excel at the challenges for which they were designed, but are often unable to handle even toy problems located within a different task domain. As a result, these systems are ill-equipped to handle unexpected situations and experiences. Commonsense reasoning can better equip these robots to handle the diverse array of "mundane" situations that occur within the real world.

Over the past fifty years, several approaches to developing synthetic commonsense have been undertaken [9]. These approaches can be roughly partitioned into two separate categories, those based on first-order logic formalisms and those with designs not originating from logic formalisms. Logic-based approaches generally utilize large collections of rules and knowledge facts and reason with logic formalisms. These approaches include Event Calculus [9], Situation Calculus [10], and action languages, including STRIPS [11], and ADL [12]. Non-logical commonsense reasoning methods include qualitative reasoning [13], the analogical processing Structure-Mapping Engine [14], probabilistic reasoning [15], and those combining diverse arrays of reasoning such as Polyscheme [16] and ThoughtTreasure [17].

Unfortunately, development and efficient utilization of commonsense reasoning is a non-trivial task that is not yet completely solved. One of the overriding difficulties is the immense collection of knowledge required before competent commonsense behavior can be achieved. It has been speculated that commonsense databases may require a few hundred million pieces of knowledge before achieving human level reasoning [3]. In order for a commonsense reasoning system to be both complete and sound, careful engineering of the facts that compose the commonsense knowledge base may be required. It is unlikely that one team can achieve this task without assistance and the overall feasibility of this approach is questionable. Compounding the problem, completeness may be impractical as no known algorithm can handle an uncertain problem context when presented with a large database of commonsense knowledge [8]. Heuristics will be required that inherently introduce unsoundness into a system.

Cyc, one of the largest commonsense databases presently available, has been developed by a small team of knowledge engineers and still possess under three million facts [18]. The team admits that it may still be one or two orders of magnitude away from the requisite amount of commonsense knowledge, despite incurring substantial development costs [19]. Other groups have attempted to leverage the Internet, by using the general population as a commonsense knowledge resource [20]. With the creation of specially designed websites, laypeople are able to assert new facts (and possibly critique existing entries). This form of human computation has several drawbacks. Without specific training, the entered facts can be too specific, too general, contradictory, in an unusable form, or simply incorrect [21].

Even without the errors from untrained Internet users, commonsense reasoning is fraught with erroneous decisions and unexpected occurrences. Humans often fail to construct appropriate inferences and frequently generate paths of reasoning that are illogical. Despite these errors, humans generally manage to solve the tasks present in daily life.

Commonsense reasoning often does not reason about all possible cases a situation can be in, but rather the default general cases. Some commonsense systems implement *Default Reasoning*, which entails assuming the default case for a particular situation and constantly changing that assumption as new information becomes available [9]. Special or exception cases are generally not handled or considered. This is in stark contrast to the goals of SA research, which aims to allow an individual to have a complete understanding of the environment, including these error cases.

Robots often operate within rich environments filled with numerous dynamic obstacles and challenges. Time and observability limitations preclude a complete understanding within real-world domains, but tasks must be completed. As robotic technology improves the generation of robot SA, robots will be better equipped to detect opportunities within the environment and circumvent unnecessary challenges. The introduction of commonsense reasoning into robotic systems will generate inferences and a level of situational understanding required to focus a robot's attention to non-obvious opportunities within the environment and handle many of the "simple" tasks that humans complete on a daily basis, but which robots fail to complete.

III. DESIGN

Robots in complex domains often require a detailed and comprehensive understanding of their environment and internal status. Minor deviations from standard operation can result in disastrous ramifications. Unfortunately, a complete understanding may be infeasible as the sheer volume of data and observability problems presented to robots can preclude perfect understanding of their environments. In these cases, heuristics, defaults, and assumptions must be used to counteract uncertainty, even if robots directly reasoning over that lack of certainty. Many means of action selection under uncertainty exist, such as selecting the choice with the highest probability or greatest utility. These decision mechanisms can all be considered heuristic, beneficial approaches when a result cannot be guaranteed. Even complete world knowledge can be insufficient if the robot lacks appropriately reasoning mechanisms as many forms of reasoning, which current robots rarely possess, are often required [3].

Commonsense reasoning may aid robots in achieving and maintaining comprehensive SA. Through default generation, commonsense reasoning may providing default actions and information, when a sensor data deficiency exists, to formally reason about the situation. In domains where SA is critical to task completion, time or other resource constraints limit the amount of reasoning and action execution that is possible before a task fails. Commonsense defaults can direct a robot's efforts towards identifying an optimal or acceptable solution in less time and with fewer consumed resources.

Several existing systems already leverage commonsense knowledge to improve performance [21], [22]. The Open Mind collection of commonsense knowledge bases has previously been used for word sense disambiguation and textual affect sensing, two areas beneficial to human-robot interaction [21]. Along with supporting Belief-Desires-Intentions type reasoning, commonsense knowledge has aided robots utilizing probabilistic approaches. Commonsense knowledge regarding the placement and availability of objects within indoor environments has previously been combined with probability theory in order to increase a robot's ability to locate objects and classify spatial locations [21], [22].

Commonsense reasoning can also aid robots detecting error cases. Often, sensors provide data that appears reasonable for a given situation even though an error case has occurred. While commonsense reasoning capabilities may only provide a listing of the default or standard status messages that a robot should be sensing, this provided data may allow robots to detect error cases and correct them. Multiple conflicting scenarios will likely result in a robot being unable to identify which are in error and which are correct, but the presence of multiple scenarios can result in the identification of a possible error situation. In this situation, the robot may employ alternate reasoning methods and activate higher cost cognitive and sensing processes to gather additional information. Results from this additional processing can reveal flawed logic on the robot's part or present missing environmental status information. This improved understanding of the world can lead the robot to identifying correct scenarios, considering new actions, and minimizing incorrect action selection.

Both SA and commonsense reasoning can be considered unsound and incomplete cognitive phenomena. Despite these failings, humans are often able to operate at high levels of efficiency within novel situations and domains. Both require defaults and expectations for humans to circumvent unobservability and make decisions with uncertainty. Additionally, both phenomena utilize diverse sensory inputs to reason over multiple mental realms [8] and mental models [2].

Sensing and expert-level reasoning over sensory inputs has improved greatly; however, robots are still unable to consistently maintain high levels of SA across multiple domains or when presented with novel situations. For example, despite being able to detect audio discontinuities, a robot must be programmed to investigate the occurrence. Motor stalling indicators may inform a robot of a locomotion error, but without the proper insight as to the cause of the problem, corrective actions cannot be taken reliably. Humans suffer from this same dilemma. Numerous avionics accidents have been traced back to flight crews that knew a piece of hardware had failed, but the crew failed to take appropriate corrective action or air traffic controllers that made disastrous decisions. In these scenarios, critical information necessary to correct the situation was either misunderstood or ignored [1], [7].

Commonsense reasoning and fact filled databases can be equipped to provide this missing logic for robotic systems. As general purpose commonsense functionality is developed, robots may be able to complete the missing functionality required to connect sensor information with data processing and beneficially influence the action selection process. Real world objects generate an almost infinite array of emergent interactions, but commonsense reasoning can minimize the importance of explicitly reasoning over those interactions. With the reduction in importance of this form of reasoning, robots may be better equipped to develop and maintain SA while reducing the system design burden. Much of the knowledge required for robotic systems to realize these benefits is not robot specific, allowing for the use of existing general purpose data repositories. The same facts that humans use to complete tasks, such as "turning on a light switch results in turning on a light", will enable robots to complete more challenging and diverse tasks. Some portions of commonsense will be domain and task specific; however, for example, knowledge of social customs may only be pertinent for tasks involving human-robot interaction.

Unfortunately, commonsense may not be amenable to direct control of robotic systems. Since commonsense primarily reasons over default situations, a commonsense based system may be ill-equipped to handle the endless series of both minor and major errors and exception cases that haunt robots. Additionally, commonsense reasoning systems generally utilize a small number of reasoning methods that may not encompass the range of expert algorithms required by robots operating in complex domains.

This paper proposes a robot design utilizing commonsense not as a top-down master process, but rather as a subordinate. In this capacity, commonsense reasoning will be continuously presented with the robot's current model of the world and self. The commonsense reasoning system will utilize facts and rules to generate lists of inferences and action suggestions for the robot. These action suggestions and inferences may refer directly to possible actions to be taken, but will also include mental operations. Proposed inferences will naturally influence the robot's data stores, but the suggested actions can include purging stale data, increased processing of a particular set of sensor data, or even shutting down poorly behaving cognitive processes.

Using commonsense reasoning in this manner requires care to prevent performance degradation below that required for real-time execution in dynamic environments. Depending on the robot's capabilities, commonsense reasoning may generate a large volume of inferences, potentially requiring substantial CPU cycles. These reasoning systems may also continuously present the same inferences and suggestions, resulting in a chronic waste of cognitive resources.

Incorporating commonsense reasoning into a robotic system as a subordinate process generates few restrictions. Purely logic-based commonsense reasoning subsystems will require some level of symbolic capabilities, but any reasoning methods, from purely logic-based to probabilistic and utility based reasoning, can be used to integrate the generated inferences and suggested action selections. The output from the commonsense reasoning subsystem acts as another sensor data stream that possesses inferences and suggestions for improving the robot's reasoning.

A. System

The utilization of commonsense as a subordinate process is a general approach that can be applied to many robotic designs and systems. To demonstrate the feasibility of this method, a proof of concept system implementation was developed that utilizes an implementation of the EM-ONE commonsense reasoning system [8] along with the CoSy Architecture Schema Toolkit (CAST) [23]. Connections to robotic sensors and actuators were achieved through the robotic hardware server Player [24]. The system was tested via simulation with the 3D high fidelity simulator Gazebo [25]. Many of the design decisions were influenced by previous research into providing robotic SA [26].

CAST is a toolkit designed to simplify the system development process by following the CoSy Architecture Schema (CAS) [23]. This modular framework allows the creation of robots capable of parallel modular processing, metacognition, and structured data management. Each loosely coupled module is a subarchitecture that generally works within one realm of reasoning. Subarchitectures consist of four component types. Unmanaged processes are light weight processes that run continuously. Working memory components store data derived from the unmanaged components and are world readable. Except for a few privileged processes, write access is only available to processes within a given subarchitecture. Managed processes, which generally perform expensive calculations, monitor working memory for relevant changes and when detected, determine what tasks should be completed. Permission is required from the task manager before a managed component can begin any new processing.

CAST was selected as the primary robotic substrate because many of its features align with the required architecture features for generating robotic SA [26]. Due to CAST's inherently modular design, data encapsulation can limit the volume of information that passes through the system. Through the use of working memories, the framework provides support for labyrinthine designs [27], as well as flexible and mutable input and output channels.

Currently, the robot exists as a collection of four CAST processes, all within one sub-architecture. Sensors are read via the unmanaged RobotSensors module that captures readings and converts them into symbolic form. The exact nature of this conversion is beyond the scope of this paper. Both raw sensor data and derived symbolic constructs are transferred to the subarchitecture's working memory. When sensor data updates are detected, the Tasker module processes the new information and incorporates it into the robot's current information data store. The Tasker provides the robot's primary decision making and action selection. As tasks are selected for execution, they are passed to the Action module, which converts the primarily symbolic based task descriptions into raw actuator signals and calls in Player. Parallel to this sensing-action selection cycle, the Tasker calls the Commonsense module. This module utilizes a symbolic representation of the current environmental state and the robot's internal status as input into the EM-ONE commonsense reasoning implementation. When the reasoning process completes, the results are passed to the Tasker. The Tasker integrates the resultant suggestions into the robot's representation of the world and self and acts on any provided suggestions or inferences. Actions suggested by the commonsense module are not immediately executed, but are reasoned over to ensure they are both beneficial to the robot's current goals and priorities and will not counteract any current or recent activities. Additionally, action suggestions are validated to ensure the actions' benefits outweigh any detrimental effects that can result from execution. Reasoning over the commonsense reasoning subsystem's outputs is required because of the incomplete and unsound properties of EM-ONE. A benefit results from this additional processing, erroneous inferences resulting from incomplete and contradictory inputs will be mitigated.

Many commonsense reasoning approaches can provide robotic systems with defaults, but EM-ONE [8] was selected because it possess properties uncommon among commonsense approaches. The most notable is the use of critics and narratives, which together provide an implementation of the critic-selector model found within the Emotion Machine [3]. Critics recognize problems, reason over previous actions, and propose solutions. Instead of a collection of facts, EM-ONE's critics utilize a collection of narratives that embody the commonsense knowledge within the system, while limiting the range of inferences that can be generated. Situations requiring SA are frequently time-constrained and brute-force or exhaustive searches within multidimensional spaces will be insufficient. Narratives, which inherently provide context to encapsulated commonsense knowledge, may reduce superfluous inferences while providing pertinent suggestions. Narratives are relatively easy to develop as the context ingrained within each narrative can limit unintended interactions and may be amenable to machine learning. EM-ONE can also be built upon alternate substrates, offering different levels of functions and capabilities. Additionally, EM-ONE provides multiple levels of thought by implementing the critic-selector model with a collection of reactive, deliberative, reflective, and meta-managerial critics. Robots aiming to develop SA must be aware of the environment's current conditions and the occurrences that led to the current situation. EM-ONE's capability for reflective thinking can enable a robot to achieve greater expressibility, while reflectively inspecting its progress to improve future decision making.

EM-ONE is a new commonsense reasoning system, but it does have limitations. Presently, few general purpose critics and narratives have been developed and validated across different domains and tasks. For our implementation, a number of new mental critics and narratives were generated for our tasks. As additional tasks and domains are explored, additional narratives and critics will be developed, and once added to the EM-ONE implementation, will generate richer inferences with greater diversity. While a given narrative or critic may be inapplicable for a particular situation or task, as the collection of critics and narratives is bolstered, the availability of applicable critics and narratives will increase. As with other existing commonsense reasoning systems, time and effort will be required to develop the system's knowledge base. The naming of EM-ONE critics follows the convention described in [8]. The general format consists of (level of thought*suspected problem=>possible solution). As an example, the critic (reactive*actor-intends-action=>do-action), see Fig. 1, received its name because it provides reactive thinking, determines if a robot intends to perform an action, and suggests the robot perform the desired action. Additionally, EM-ONE lacks general capabilities for handling the commonsense law of inertia and forms of indirect effects [9].

```
(defcritic
 (reactive*actor-intends-action=>do-action)
 (in conditions current-conditions
  (intends :actor ACTOR :prop (ACTION
                        :actor ACTOR :object OBJECT)))
 (=>)
 (in conditions current-conditions
   (assert (does :actor ACTOR :prop (ACTION
                       :actor ACTOR :object OBJECT)))))
```

Fig. 1. Example of a EM-ONE critic used during the evaluation.

IV. Evaluation

A simulated robot task exploring a disaster response scenario was developed to provide a proof of concept. For this experiment, 5 new critics and 10 new narratives where combined with 2 existing critics. Additional critics and narratives described in [8] were not applicable and were not used. The robot was directed to acquire video data from the interior of a building, which had received damage from an explosion, so that first responders could safely determine if the building was structurally sound. The responders supervising the robot assumed that the building was devoid of people, animals, and hazardous objects. Instead of an empty building, the robot, via exploration, discovers a first responder, a wounded victim, a wild animal, and a suspicious looking package surrounded by debris. While the robot possessed the capabilities to handle these unexpected occurrences, it was not explicitly programmed to identify or react to them.

The computational load incurred during the initial evaluation period for sensor readings was reduced by approximated functional simulations. Sensing capabilities included the ability to detect: animals, humans, medium sized objects, debris, and colored substances. Simulating the presence of RFID markers, humans were identified as either a first responder (if in close proximity to an RFID) or a victim. Continuous distances to objects were sensed, but distances were partitioned into categories: far, near, and close. The robot could sense whether two objects were either touching or very close together.

A. Identifying a Wounded Victim

When the robot encountered the victim located within the building, the RobotSensors module detected a human without an RFID within close proximity, who was touching a red substance. Except for obstacle avoidance, the robot's standard programming would not have interacted with the victim. The commonsense module created several inferences and suggested courses of action. The EM-ONE critic (reactive*object-relation-infers-property=>classify-object) inferred that the human was an injured victim. A second critic, (reactive*infer-problem=>suggest-action), inferred that the victim required help. Additionally, the critic suggested that the robot ask the victim if assistance was required. The tasker developed a plan to interact with the victim. Inferences were generated once the victim became visible, but the tasker accounted for the distance from the victim and waited to begin interaction until the robot's speaker was close enough to be heard. Fig. 2 shows the narrative used to infer that the victim was hurt.

```
(defnarrative victim-is-hurt
  (isa :x human :class victim)
  (is-touching :subject red-substance
               :object human [1])
  (isa :x human :class hurt-victim [2])
  (implies [1] [2]))
```

Fig. 2. Example of a EM-ONE narrative used during the evaluation.

B. Avoiding Animals

When the robot approached an animal in the building, the tasker took no actions. Once the robot was near the animal, the EM-ONE critic (reactive*see-potential-problem=>suggest-action) engaged. Using the narrative see-animal-leave, the commonsense module provided the inference that the animal may be dangerous and suggested the robot should report the animal's existence and leave the area.

C. Adjusting Speed Due to Debris

While traveling to the area where the suspicious package was located, the robot found some debris. Using the critic (reactive*see-potential-problem=>suggest-action2) and the two separate narratives see-debris-go-slow and see-debris-go-slow2, inferences that the debris may be dangerous and moving slowly was prudent were generated.

D. Identifying Packages

While traveling through the area covered with debris, the robot observed a package covered in a white substance. Similar to the victim encounter, the (reactive*infer-problem=>suggest-action) engaged. Since different narratives were used by the critic, the resultant inferences and action suggestions differed from the victim discovery. The new inferences were that the package was both a suspicious-object and a dangerous-object. The commonsense module suggested that the robot examine the object in order to determine the object's importance.

E. Interacting with First Responders

After the above events, an unexpected first responder was discovered. Initially, the responder's RFID was not sensed and the inference that the human was a responder was not made and no EM-ONE critics engaged. After traveling closer to the human, the RFID was sensed and the human was classified as a first responder. Upon identification, the EM-ONE critic (reactive*see-object-existing-object=>take-action) engaged. This critic utilized three separate narratives, describe-animal, describe-victim, and describe-package, to generate suggestions that the robot report its previous findings to the first responder. The tasker permitted these suggestions and the information was transfered to the first responder.

F. Discussion

Although contrived, the situations appearing in the proof of concept system demonstrated the potential for increased system performance when commonsense acts as a subordinate process. Even though the robotic system was equipped with the basic functionality required to observe and act upon several situations located within the test environment, the non-commonsense portion of the robotic system was not equipped to understand or shape the action selection process. The commonsense module's introduction allowed the system to successfully diagnose the situation and present the tasker with the required information to correctly handle the obstacles encountered by the robot. For example, in Sections IV-D and IV-E, several critics were engaged multiple times. On each occurrence, a different narrative was used, demonstrating the general applicability of EM-ONE's critics.

The presented approach can be integrated into many existing systems, providing increased reasoning capabilities while causing minimal disruption. Inferences and defaults generated by commonsense reasoning can alter symbolic reasoning within a robot or modify a robot's belief states and utilities within probabilistic systems. Since the generated commonsense inferences are reasoned over, errors resulting from unsoundness will be minimized and missing inferences due to incompleteness will not have a negative affect.

V. CONCLUSIONS AND FUTURE WORK

Robots will continually be tasked with new and challenging domains and goals. Improving robot SA has the potential to assist in these endeavors and increasing commonsense levels in robotic systems is a viable approach to achieving high levels of robotic SA.

This paper presented a general method for utilizing commonsense reasoning to improve robot SA. Instead of directly driving and manipulating the system, commonsense reasoning provided suggestions and inferences to be reasoned over by the system. By allowing robots to reason over default inferences, negative effects of commonsense's incompleteness and unsoundness can be minimized.

An initial proof of concept system was developed demonstrating a robot tasked with exploring an indoor environment. During the task, the robot encountered several situations where explicit programming was lacking. Through the use of an integrated commonsense module, appropriate actions were inferred and subsequently taken by the robot.

The presented work is part of a larger effort exploring the effectiveness of different architectural features and design mechanisms required by robots to achieve and maintain SA.

REFERENCES

[1] M. J. Adams, Y. J. Tenney, and R. W. Pew, "Situation awareness and the cognitive management of complex systems," in *Human Factors*, vol. 37, no. 1, 1995, pp. 85–104.

[2] J. Uhlarik and D. A. Comerford, "A review of situation awareness literature relevant to pilot surveillance functions," Office of Aerospace Medicine, Federal Aviation Administration, U.S. Dept. of Transportation, Washington, DC, Tech. Rep. DOT/FAA/AM-02/3, 2002.

[3] M. Minsky, *The Emotion Machine*. Simon & Schuster, 2006.

[4] J. M. Flach, "Situation awareness: Proceed with caution," in *Human Factors*, vol. 37, no. 1, 1995, pp. 149–157.

[5] N. B. Sarter and D. D. Woods, "Situation awareness: A critical but ill-defined phenomenon," in *The Int. J. of Aviation Psychology*, vol. 1, no. 1, 1991, pp. 45–57.

[6] M. R. Endsley, "Situation awareness global assessment technique (SAGAT)," in *Proc. Nat. Aerospace and Electronics Conf.*, 1988, pp. 789–795.

[7] ——, "Toward a theory of situation awareness in dynamic systems," in *Human Factors*, vol. 37, no. 1, 1995, pp. 32–64.

[8] P. Singh, "EM-ONE: An architecture for reflective commonsense thinking," Ph.D. diss., Massachusetts Institute of Technology, 2005.

[9] E. T. Mueller, *Commonsense Reasoning*. Morgan Kaufman, 2006.

[10] R. Reiter, *Knowledge in action: Logical foundations for specifying and implementing dynamical systems*. MIT Press, 2001.

[11] R. E. Fikes and N. J. Nilsson, "STRIPS: A new approach to the application of theorem proving to solve problems," *Artificial Intelligence*, vol. 2, pp. 189–208, 1971.

[12] E. P. D. Pednault, "ADL: Exploring the middle ground between strips and the situation calculus," in *Proc. First Int. Conf. on Principles of Knowledge Representation and Reasoning*, R. J. Brachman, H. J. Levesque, and R. Reiter, Eds. San Mateo, CA: Morgan Kaufmann, 1989, pp. 324–332.

[13] J. de Kleer, "Qualitative and quantitative knowledge in classical mechanics," Artificial Intelligence Laboratory, Massachusetts Institute of Technology, Tech. Rep. AITR-352, 1975.

[14] B. C. Falkenhainer, K. D. Forbus, and D. Gentner, "The structure-mapping engine: Algorithm and examples," *Artificial Intelligence*, vol. 41(1), pp. 1–63, 1989.

[15] S. Thrun, W. Burgard, and D. Fox, *Probabilistic Robotics*. Cambridge, MA.: The MIT Press, 2005.

[16] N. L. Cassimatis, "Polyscheme: A cognitive architecture for integrating multiple representation and inference schemes," Ph.D. diss., Massachusetts Institute of Technology, Cambridge, MA, Feb. 2002.

[17] E. T. Mueller, *Natural language processing with ThoughtTreasure*. New York: Signiform, 1998.

[18] "ResearchCyc," Website, Sep. 12 2008. [Online]. Available: http://research.cyc.com

[19] P. Singh, "The open mind common sense project," KurzweilAI.net, 2002, downloaded January 15 2008. [Online]. Available: http://www.kurzweilai.net/meme/frame.html?main=/articles/art0371.html

[20] P. Singh and B. Barry, "Collecting commonsense experiences," in *Proc. 2nd Int. Conf. on Knowledge Capture*, 23-25 2003.

[21] R. Gupta and M. J. Kochenderfer, "Common sense data acquisition for indoor mobile robots," in *Proc. 19th Nat. Conf. on Artificial Intelligence*, 2004, pp. 605–610.

[22] T. Kollar and N. Roy, "Utilizing context when planning to find things," in *Proc. Int. Conf. on Robotics and Automation*, 2009, In Press.

[23] N. Hawes, A. Sloman, J. Wyatt, M. Zillich, H. Jacobsson, G.-J. M. Kruijff, M. Brenner, G. Berginc, and D. Skocaj, "Towards an integrated robot with multiple cognitive functions," in *Proc. 22nd Conf. on Artificial Intelligence*, AAAI Press. Menlo Park, California: AAAI Press, 2007, pp. 1548–53.

[24] B. P. Gerkey, R. T. Vaughan, and A. Howard, "The Player/Stage Project: Tools for multi-robot and distributed sensor systems," in *Proc. Int. Conf. on Advanced Robotics*, Portugal, 2003.

[25] N. Koenig and A. Howard, "Design and use paradigms for gazebo, an open-source multi-robot simulator," in *2004 IEEE/RSJ Int. Conf. on Intelligent Robots and Systems*, Sendai, Japan, 2004, pp. 2149–2154.

[26] S. T. Freedman and J. A. Adams, "Synthetic cognitive agent situational awareness components," in *AAAI Tech. Report FS-08-04*, 2008.

[27] A. Sloman, "On designing a visual system," *J. of Experimental and Theoretical Artificial Intelligence*, pp. 289–337, 1989.