

Heterogeneous Artificial Agents for Triage Nurse Assistance

D.M. Wilkes, *Member, IEEE*, Stan Franklin, Erdem Erdemir, *Member*, Stephen Gordon, *Member, IEEE*, Steve Strain, Karen Miller and Kazuhiko Kawamura, *Fellow, IEEE*

Abstract— A dream of humanoid robot researchers is to develop a complete “human-like” (whatever that means) artificial agent both in terms of body and brain. We now have seen an increasing number of humanoid robots (such as Honda’s ASIMO, Aldebaran’s Nao and many others). These, however, display only a limited number of cognitive skills in terms of perception, learning and decision-making. On the other hand, brain research has begun to produce computational models such as LIDA. In this paper, we propose an intermediate approach for body-brain integration in a form of a scenario-based distributed system. Busy hospital Emergency Departments (ED) are concerned with shortening the waiting times of patients, with relieving overburdened triage team physicians, nurses and medics, and with reducing the number of mistakes. Here we propose a system of cognitive robots and a supervisor, dubbed the TriageBot System that would gather both logistical and medical information, as well as take diagnostic measurements, from an incoming patient for later use by the triage team. TriageBot would also give tentative, possible diagnoses to the triage nurse, along with recommendations for non-physician care. Some of the robots in the TriageBot System would be humanoid in form, but it is not necessary that all of them take this form. Advances in humanoid robotic design, in sensor technology, and in cognitive control architectures make such a system feasible, at least in principle.

I. INTRODUCTION

When we think of an emergency department (ED) we often think of severe trauma patients arriving by ambulance or even by helicopter. However, there are many patients with significantly less severe ailments, who arrive by car or walk in on foot. For these patients, the triage nurse provides a vital role by performing tasks such as gathering data from

the patient, taking diagnostic measurements, assessing the severity of the patient’s condition for ordering priority of treatment, and updating the patient’s data at timely intervals. While there may be a role for robots to play in the severe trauma situations, this paper focuses on the less severe cases with particular emphasis on the role of the triage nurse and how robots may be used to assist the nurse in the performance of his/her duties. There have been other robotic projects designed to assist in health care, such as the Nursebot Project which assists the elderly [1], but the system of this paper is the first to address the use of robots in the ED.

Modern emergency triage requires safe and efficient operations to deliver acute care. Triage is simply the allocation of emergency services to patients based on the severity of their condition when resources are scarce. The intent is that the sickest patient receives emergency care first [2]. Modern emergency triage in acute care is no longer a simple act of sorting ED patients to prioritize the sickest patient to have emergency care. Current overcrowding in hospitals demands ED’s to hold patients waiting for admission. This causes a back up in ED triage and ED throughput to a much slower pace [3]. To improve ED throughput triage teams accomplish an extensive assessment of vital signs, electrocardiogram, laboratory tests, medical history and social history (smoking, drug, home medications, years of education, etc.). Often triage starts treatment to include the administration of intravenous fluid, medications, splinting, bandaging and ordering other tests [4][5]. Triage clinicians are further responsible to enter this information into the electronic medical record to communicate to the emergency department core and hospital [6]. Triage is now the intake center of the ED and often the hospital [7]. The numbers of rooms, personnel, and different disciplines required for triage depends on the number of patient visits per day at an individual emergency department. Management of this complex system is critical to the fair distribution of emergency services. Coordinated by a registered nurse, the triage team is comprised of registration clerks, paramedics, nurses, physicians and police. There is inherent variability in patient arrival times. Timing of acute illness and injury is difficult to predict as is the severity of the presentations [8][9]. This unpredictability plus overcrowding, limited clinical personnel and healthcare resources can quickly lead to a hectic situation with long waiting times and poor healthcare consequences. [9][10].

The Institute of Medicine has identified ED overcrowding as a major public health problem [11]. Increases in patient presentations and the ED boarding of hospital admitted patients waiting for hospital beds cause congestion in the triage area and prolonged waiting times for treatment. It is

This work was supported in part by under NSF grants EIA0325641, “ITR: A Biologically Inspired Adaptive Working Memory System for Efficient Robot Control and Learning”.

D.M. Wilkes is an Associate Director of the Center for Intelligent Systems and Associate Professor of Electrical and Computer Engineering at Vanderbilt University, 37235-0131 (email: mitch.wilkes@vanderbilt.edu). Stan Franklin is the W. Harry Feinstone Interdisciplinary Research Professor, Cognitive Computing Research Group, Institute for Intelligent Systems, The University of Memphis, Memphis, TN 38152, (email: franklin.stan@gmail.com). Erdem Erdemir is a Research Assistant in the Department of Electrical Engineering and Computer Science, Vanderbilt University, (email: erdem.erdemir@vanderbilt.edu). S.M. Gordon graduated in 2009 with a Ph.D. in Electrical Engineering from Vanderbilt University, (email: stephenmgordon@comcast.net). Steve Strain is a graduate student in the Cognitive Computing Research Group, Institute for Intelligent Systems, The University of Memphis (email: sfstrain@yahoo.com). Karen Miller is Director of Clinical Research Operations for Emergency Medicine at Vanderbilt University, (email: Karen.f.miller@vanderbilt.edu). K. Kawamura is the Director of the Center for Intelligent Systems and Professor of Electrical and Computer Engineering at Vanderbilt University, Nashville TN 37235-0131 (e-mail: kawamura@vuse.vanderbilt.edu).

reported that prolonged ED waiting times have poor outcomes. Increased patient mortality, time to treatment for infections, blood clots and pain are reported. Patient satisfaction decreases with long waiting times [10]. There are anecdotal and media reports of deaths in the ED waiting room [12]. There are health care disparities reported regarding increased waiting times to see the provider in the United States for the uninsured, low-income, African-American and Hispanic populations and female gender [13]. There are further implications that a crowded ED negatively impacts the ED's ability to respond in a mass casualty situation [11]. In contrast, team triage systems, computerized triage adjuncts and placing physician orders at triage improves ED throughput [5]. Novel strategies for ED triage are necessary to meet the resource and patient-safety demands of the acute care emergency setting. Robot assistants in the ED triage may improve ED throughput and provide a safer environment.

Let us review a typical scenario involving the interaction between a patient, a nurse and other personnel in the ED. Patients present to the ED by ambulance, ambulatory or wheelchair assist. They are alone or with one or more family and/or friends. ED patient visitors arrive at any time before or after the patient. A police or security officer is in or around the entry points to triage. When the patient arrives they must register. During this registration process a variety of information is gathered, such as patient's primary complaint and data (including age, gender, race and primary language) as well as consent forms filled out and signed. Typically this is accomplished by a registration clerk. If the patient's primary complaint seems life threatening (chest pain, profuse bleeding, loss of consciousness or difficulty breathing), the registration clerk alerts the clinical staff for immediate evaluation. If not, the registration process is completed and the patient either waits in the waiting room or is evaluated by the clinical staff. The clinical staff collects basic diagnostic data is measured, including blood pressure (BP), pulse rate, blood oxygen saturation, respiration rate, height and weight. Additionally, the nurse asks the patient questions including: "What is the chief complaint?" and "Where is the pain?". Most hospitals' ED triage intake information has grown to include additional assessments including highest level of education, vaccine history and allergies. Much of this information has become required computer fields for the nurse to complete before the patient may move through the system [6]. The nurse may also use Visual Analog Scores to assess pain levels, shortness of breath, etc. These Visual Analog Scores are graphical icons used to depict severity of pain and discomfort. From this information the nurse produces an Emergency Severity Index (ESI) score [14], assessing the patient's condition. An additional ESI score may be estimated by a computer program as well [15]. At this point, the clinical staff makes a schedule to check back with the patient on a timely basis, for example every 10 minutes, where the time interval is a function of the patient's condition. Finally, the patient is assigned an ED bed or goes to the waiting room of the ED where the triage team is responsible to observe the patient until assignment to an ED bed. An illustration of this

scenario for an ED that has 100-125 visits per day is shown in Figure 1a.

A safer more efficient system using robot assistants is proposed in Figure 1b. In this scenario, robots are placed to assist the triage team. The triage team and patients interact with the robots to register and make initial assessments. The robots are able to input into the electronic medical record in real time. Robots help patients or family enter their own information using a touch-pad computer screen to relieve the majority of the data entry burden from the triage team. The robots update the patients on current wait times. The triage teams use the robots to, through the use of cameras; visually inspect and listen to the waiting room and/or a specific patient. The robots could relay pre-programmed critical information to alert the clinicians in case of a situation that requires immediate clinical attention like chest pain. The clinical teams use information gathered by the robot freeing the triage teams for patient care and to determine the best care plan for each patient. Use of the triagebot system improves efficiencies, throughput and patient safety.

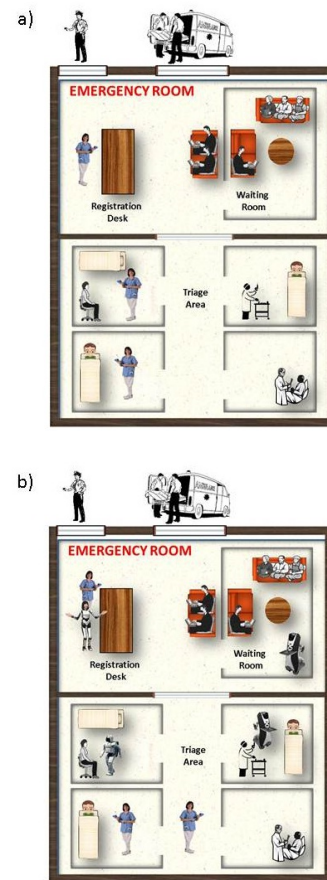


Figure 1: An emergency room environment, (a) current environment with not robots and (b) with robot triage nurse assistants.

II. SYSTEM OVERVIEW

The TriageBot System consists of an overall supervisor responsible for the "big picture" of all the activities in the ED and a collection of individual robots responsible for

various activities within the ED. We envision the supervisor as well as each of the robots being implemented as cognitive agents, and these cognitive agents would then coordinate and collaborate with each other. While it is also the case that the broader role of the supervisor would entail it giving some commands to the robots, it will not be a master-slave system in which the robots are somewhat dumb and the intelligence is concentrated in the supervisor. All of the agents within the system will also communicate and interact with the doctors and nurses, and the robots will also interact with the patients. This is illustrated in Figure 2.

The robots in the TriageBot System may take many forms. Robots that need to be able to manipulate items or interact with patients using gestural interfaces may need to have a humanoid form. Robots used primarily to gather patient data, such as during the registration process may take the shape of a smart kiosk, similar to an Automatic Teller Machine (ATM) where buttons, touchscreens, digital signature pads and possibly voice (e.g., as in the Dragon Medical Mobile Search Application by Nuance) are well suited to gathering information. Other robots may need to be mobile in order to transport various items. A possible platform might be the PR2 robot from Willow Garage. A robot designed primarily to measure diagnostic data may have the form of a “smart chair” equipped with sensors for measuring blood pressure, blood oxygen saturation, pulse rate, respiration rate, height and weight. However, in all cases we envision the robot to need a cognitive architecture to provide the combination of intelligence and adaptability to deal with the dynamic complex requirements of the ED. Additionally, they will need network connections to provide full communication with the supervisor, doctors and nurses. The network is most likely to be a typical Local Area Network (LAN) and communication between the elements supported by appropriate middleware such as YARP described in [16] [17].

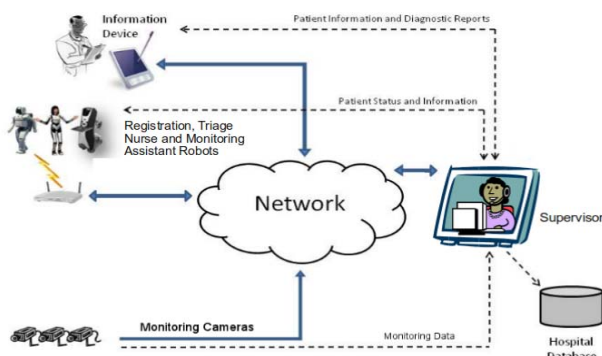


Figure 2: A network diagram for the TriageBot System.

III. THE MAIN ELEMENTS OF THE TRIAGEBOT SYSTEM

The main robots and agents of the TriageBot System are now described in greater detail. We will describe them in the order in which a patient is likely to encounter them.

A. Robot Registration Assistant

The Robot Registration Assistant would be at the registration desk and would be the first robot encountered by the patient. This robot will have a humanoid form and will need a high degree of ability to interact with the patient. It must be able to recognize humans and track them in its environment. Additionally, it must be able to engage in basic conversation with the registering patients. Of course, it must also be able to interact with the other agents in the system.

This robot will gather basic patient data such as name, address, telephone numbers, insurance information, etc. It will also start gathering some diagnostic data by asking such questions as “What is the chief complaint?”, “Where is the pain?” and “What is the level of pain?”. Visual Analog Scores may be used to assess pain levels, shortness of breath, etc. The methods of interaction used for gathering this data may include voice dialog and touch sensitive screens as may be encountered in a smart kiosk. This data is entered into the patient’s file, and then the patient is directed to the Robot Triage Nurse Assistant for gathering other diagnostic data.

B. Robot Triage Nurse Assistant

This robot is likely to have a specialized form designed specifically for taking measurements. A likely form is that of a chair instrumented with the necessary sensors. These measurements will include blood pressure, pulse rate, blood oxygen saturation, respiration rate, height and weight. From this information an ESI score, assessing the patient’s condition and priority in the triage queue, may be calculated and all data entered into the patient’s file where it is available to the other agents in the system including the doctors, nurses and medics.

In general, the high-level interaction skills of this robot are less complex than the others, since its duties are more specifically prescribed. On the other hand, it will require a higher level of motor skills since it will be responsible for taking measurements directly from the patients. After this data is gathered, the patient is sent back to the waiting room where he/she will be monitored while waiting for treatment.

C. Robot Monitoring Assistant

After reviewing all the data collected, the Robot Monitoring Assistant selects an appropriate time interval for checking up on the patient in the waiting room. This robot will periodically check to see if the patient is still in the waiting room, if they are conscious, and possibly take simple measurements such as blood pressure and pulse rate. Additionally, it may inquire about the level of pain. There is some flexibility in the form of this robot. It is likely to be a mobile robot and may or may not have humanoid characteristics. It will require a substantial level of cognitive skill in order to interpret and respond to a wide variety of events and interactions in the waiting room.

D. Supervisor

The Supervisor will act as the central manager of all the robots, as well as providing an interface to hospital

personnel and databases, except for the doctors and nurses that interact directly with the patients. They, of course, will still have the direct interfaces that they usually use. Additionally, there are likely to be sensors, such as cameras, monitoring the waiting room and possibly the treatment rooms. These would enable the Supervisor to check for important events including whether a patient has fallen to the floor or whether a patient is still conscious. Finally, the Supervisor may calculate possible diagnoses and suggest early testing or other non-physician care.

IV. ARCHITECTURAL CONCERNS

A. Requirements for a Cognitive Robot in a Partially Structured Environment

Humans process sensory data, and select and begin executing a response five to ten times a second [18]. Humanoid robots, operating in human-like environments should also be able to process sensory data and choose actions at a similar rate of 5-10 Hz. Humans deal continually with tremendous amounts of sensory data, much of it irrelevant, by employing their attention mechanisms as a filter. A humanoid robot “living” in a typical partially structured environment should also filter large amounts of sensory data using an attention mechanism. This implies that the robot must be capable of attentional learning, that is, of learning what to pay attention to. Such learning would seem to require both top-down and bottom-up processing, as well as the self-organization of concepts. The latter will also require self-derived representation, that is, perceptual learning. All this entails considerable bottom-up modifying of representations and organizing, combined with top-down analysis of performance.

If a cognitive humanoid robot has humans or databases readily available, say for example, via a wireless internet connection, it might not have to be widely knowledgeable, being able to ask about what it doesn’t know. That, of course, requires that it be smart enough to know when it doesn’t know [19]. In order for a cognitive humanoid robot to rely on humans or databases for knowledge, it must have enough metacognitive ability to recognize its lack of knowledge. A cognitive humanoid robot operating well in a human-like environment had best be controlled by a cognitive architecture capable of perceptual and attentional learning, as well as of higher-level cognitive processes such as metacognition [20] [21].

B. Establish dynamic sensory-behavior linkages

As the need for cognitive humanoid robots to become useful partners in our society increases, it is important to look beyond engineering-based control and learning approaches. For example, humans have the capacity to receive and process enormous amount of sensory information from the environment, integrating complex sensory-motor associations as early as two years old [22] [23]. Most goal-oriented robots currently perform only those or similar behavioral tasks they were intended for. Very little adaptability in behavior generation is exhibited when an important environmental event occurs. What is needed here

is an alternative paradigm for behavioral task learning and execution. Specifically, we see cognitive flexibility and adaptability in decision making in our brain as a desirable design goal for the next generation of cognitive robots. For example, human decision making is strongly influenced by our internal states such as emotions. A change in internal state results in changes in our perception of which goals are more important. This type of decision making leads to more “acceptable” solutions rather than precise engineering solutions.

Engineers have long used mathematical models and feedback loops to control mechanical systems. Limitations of model-based control led to a generation of intelligent control techniques such as fuzzy control, neuron-computing and reconfigurable control. The human brain, on the other hand, is known to process a variety of stimuli in parallel, to ignore stimuli non-critical to the task in hand, and to learn new tasks with minimum assistance. This process, known as cognitive or executive control, is unique to humans and some animals [24] [25] [26]. We consider this cognitive control capability as an important design principle for cognitive humanoid robots [27] [28].

C. Cognitive control architecture, perception, attention, and situational awareness

As pointed out earlier, each of the robot assistants, as well as the software agent supervisor, will require a cognitive control architecture. Perception systems will play a critical role for the performance of the motor actions in each of the robots and in the supervisor. Each robot may encounter many percepts at any given moment, and many of them may be distractors for the current task. The limited capacity property of an attention system provides focus for the robots to search for appropriate actions in order to accomplish the given tasks. A significant role of the attention system is the determination of which chunks¹ of information should be actively retained, and which may be safely discarded, for the current critical task success. Furthermore, the emergency department domain will require of each robot assistant, and of the supervisor, considerable situational awareness, that is “... the perception of elements in the environment within a volume or time and space, the comprehension of their meaning, and projection of their states in the near future.” [29] Situation awareness by triage robot assistants in an emergency department setting includes being aware of unexpected events and of the unpredictable behavior of patients. For these responsibilities, and more, we intend to use the Triage Nurse Assistant Architecture along with the LIDA cognitive control architecture derived from the LIDA cognitive model.

D. Triage Nurse Assistant Architecture

The software architecture for the triage nurse assistant robots is shown in Figure 3. The same architecture may also be used for the Registration and Monitoring Assistant Robots as well. In this architecture an array of *Perceptual*

¹ In this context, the term “chunks” is used to refer to the memory items that are utilized by the working memory.

Agents are used to detect external stimuli. These agents are designed to operate in parallel, independently perceiving information from incoming sensory data. Typically, as perceptual information is detected, that information should be sent concurrently along three separate control paths that provide for reactive, routine, and deliberative control processes.

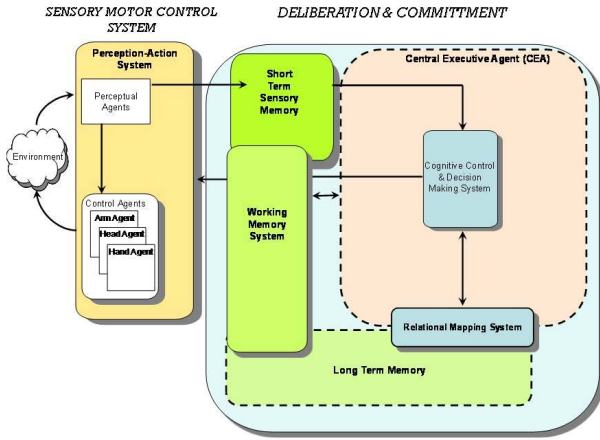


Figure 3: Triage Nurse Assistant Architecture

Perceptual information flows to the Working Memory System (WMS), where it is combined with information coming from long-term memory (LTM) and the system's evaluation and response systems. This combination facilitates categorization and compression of the incoming perceptual signals. The output of the WMS to the Central Executive Agent (CEA) is therefore a labeled and compressed representation of the current perceptual state with respect to the system's current knowledge (LTM) and goals.

As the flow of information enters the CEA, it may be used in a variety of ways. First, it may be used to trigger a new deliberative cycle, in which the system attempts to formulate a plan that helps the system meet its goals, given its current knowledge and the perceptual state. Second, the incoming information may be used to interrupt an on-going deliberative cycle. This can happen if a desired action can no longer be performed, or if an element of the compressed perceptual representation signals a high-priority state for which a plan must be developed. Finally, the incoming perceptual information may also be folded into a currently forming plan, but only if the new state information is consistent with the partial plan that has already been formed. The functions of the CEA are supported by the Relational Mapping System.

E. The LIDA Model and its Architecture

The LIDA model [20] [30] [31] is a comprehensive, conceptual and computational model covering a large portion of human cognition². Based primarily on Global Workspace theory [32] [33], the model implements and

fleshes out a number of psychological and neuropsychological theories. The LIDA computational cognitive architecture is derived from the LIDA cognitive model. The LIDA model and its ensuing architecture are grounded in the LIDA cognitive cycle. Every autonomous agent [34], be it human, animal, or artificial, must frequently sample (sense) its environment and select an appropriate response (action). More sophisticated agents, such as humans and cognitive robots, process (make sense of) the input from such sampling (be situation aware) in order to facilitate their decision making. The agent's "life" can be viewed as consisting of a continual sequence of these cognitive cycles.

F. The LIDA Cognitive Cycle

The LIDA model hypothesizes a rich inner structure of the LIDA cognitive cycle. (Please see Figure 4.) Detailed descriptions are available elsewhere [18] [35].

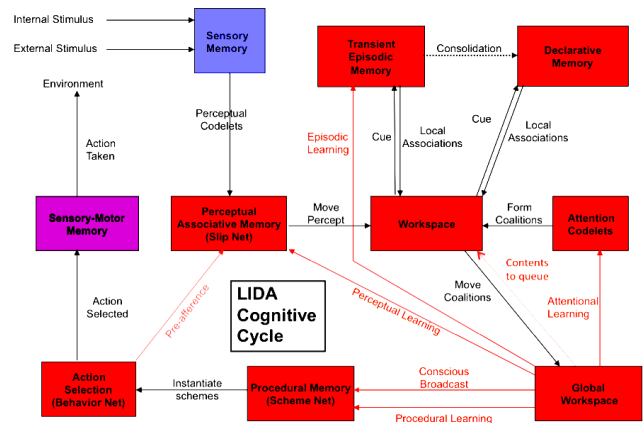


Figure 4. The LIDA Cognitive Cycle

During each cognitive cycle the LIDA agent first makes sense of its current situation as best as it can by updating its representation of its current situation, both external and internal. By a competitive process, as specified by Global Workspace Theory, it then decides what portion of the represented situation is most in need of attention. Broadcasting this portion, the current contents of consciousness³, enables the agent to choose an appropriate action and execute it, completing the cycle. Thus, the LIDA cognitive cycle can be subdivided into three phases, the understanding phase, the attention (situation awareness) phase, and the action selection phase. A cognitive cycle can be thought of as a moment of cognition -- a cognitive "moment."

A LIDA controlled robot assistant or supervisor would be capable of learning in several different modalities, perceptual, episodic, procedural and attentional. In each of these modalities, it would also be capable of instructionist

² "Cognition" is used here in a particularly broad sense, so as to include perception, feelings and emotions.

³ Here "consciousness" refers to functional consciousness (Franklin 2003). We take no position on the need for, or possibility of, phenomenal consciousness.

learning, the learning of new representations, as well as of selectional (reinforcement) learning that modifies the strength of existing representations. Such learning may well prove critical to a fully functioning TriageBot system.

V. CHALLENGES

Operating a cognitive system in the complex dynamic environment of an ED poses many challenges. Some challenges we have identified include the following.

1. Roboethics is increasingly important to the design of robotic systems, especially those in which there will be a large amount of human-robot interaction (HRI). Clearly, the proposed triage robots fall into this category. Some of these issues are discussed and linked at <http://www.roboethics.org>. An interesting discussion may be found in [36]. At all times, the health and safety of the patient must be safeguarded, thus the design of the system must take into account the relative capabilities of the robots and the human medical staff. It is important to model and design the interactions to protect the patient. Another major issue in a medical application is protection of the patient's privacy, in particular making sure to protect the medical data. The system must be designed to adhere to existing protocols for protecting patient privacy. These are just two ethical considerations among many that must be addressed.
2. The system should support Natural Language Processing in this domain. Two of our robot assistants must communicate with patients in natural language, while all, including the supervisor must so communicate with humans. LIDA's predecessor, IDA, whose task was finding and negotiating new jobs for sailors at the end of their current tour of duty [37], successfully automated all the tasks of Navy personnel officers. IDA did so using email in unstructured English [38]. Much the same technique should work for the various TriageBot robots and the supervisor, except that the vocabulary of the emergency medical domain is much more diverse and specialized, presenting a challenge.
3. Controlled by the LIDA architecture, the supervisor should be able to perform an agent-based differential diagnosis and suggest further tests to be performed. Differential diagnosis is a generic process for formulating diagnostic possibilities from the initial medical data. Conventional approaches to computer-assisted diagnosis have had only limited successes in improvement of practitioner performance, and to date have not proven effective in improving patient outcomes [39]. We suggest that an approach implementing human-style reasoning may prove successful where other approaches have failed. LIDA is considered a suitable architecture for the implementation of a diagnostic agent, and the pre-ordering of tests based

on diagnostic possibilities is an identified area for improving the workflow in a triage department [40].

4. Measuring the patient's vital signs poses several potential difficulties. The measurement sensors must be properly applied to the patient for reliable measurements. While these measurements are being taken, the system should monitor for events of major importance, such as cardiac arrest. Additionally, since the patients may be in considerable discomfort their behavior may become erratic; thus the system should monitor for this. We currently have an electrical engineering senior design project at Vanderbilt aimed at designing the robots at the registration desk and waiting room of the ED. Their primary responsibilities will include gathering patient information, some initial vital signs, and basic patient monitoring.
5. In monitoring the waiting room, the system needs to be able to interpret events such as fainting, falling or erratic patient behavior. This requires scene analysis and interpretation. Scene interpretation requires coordinating many separate tasks including occlusion reasoning, surface orientation estimation, object recognition, and scene categorization [41]. Such interpretation can be an exhaustive open-ended research project in itself, but the system can benefit from knowledge of the patient's medical record. Existing conditions or prior predisposition to fainting can lead the system to pay particular attention to specific patients in the waiting room. Additionally, it may be possible to give the patient some type of badge to wear that facilitates visual tracking. RFID tracking may also be considered. The ED domain will in particular require sophisticated event-based representations [42], a considerable challenge in an often-chaotic ED waiting room.

Triage started as a clinical function, conducted by a nurse to sort and prioritize patient urgent health needs and has become a location for ED intake with operational challenges met only by novel approaches. The function of triage has expanded to include diagnostics, treatments and assessments that take limited resources of time and healthcare personnel to accomplish. The available computer technology has stream-lined the information and has left data entry to the clinical staff [43]. Now registration clerks, paramedics, nurses, physicians, and police are members of a triage team. Most triage systems continue to depend on an experienced nurse to perform the triage task and coordinate the ED intake. This requires situational awareness of the patients in the waiting room, number of patients waiting in the ED for admission, ambulance traffic and hospital capacity. Overcrowding, limited experienced personnel and healthcare demands can quickly lead to a bottle-neck situation in the entire ED system leading to serious mistakes with untoward health consequences. A variety of strategies for triage assessment and decision making are utilized to met the

resource and operational demands of the process. [6] [8]. The innovation of healthcare robotics addresses the operational challenges of healthcare intake demands.

6. CONCLUSIONS

We have addressed some of the problems of busy hospital Emergency Departments related to shortening the waiting times of patients, relieving overburdened triage team physicians, nurses and medics, and with reducing the number of mistakes. We have proposed a system of cognitive robots and a supervisor, dubbed the TriageBot System that would gather both logistical and medical information, as well as take diagnostic measurements, from an incoming patient for later use by the triage team. Recent advances in humanoid robotic design, in sensor technology, and in cognitive control architectures make such a system feasible, at least in principle.

ACKNOWLEDGMENT

The authors would like to thank Flo Wahidi and members of the Center for Intelligent Systems for their contributions and support. They are also indebted to the members of the Cognitive Computing Research Group at the University of Memphis.

REFERENCES

- [1] Pollack, M. E., Engberg, S., Matthews, J. T., Thrun, S., Brown, L., Colbry, D., Orosz, C., Peintner, B., Ramakrishnan, S., Dunbar-Jacob, J., McCarthy, C., Montemerlo, M., Pineau, J., and Roy, N., "Pearl: A Mobile Robotic Assistant for the Elderly," *AAAI Workshop on Automation as Eldercare*, Aug., 2002.
- [2] Iserson, K. and Moskop, J., "Triage in Medicine, Part I: Concept, History, and Types," *Annals of Emergency Medicine*, Vol. 49, No. 3: March 2007.
- [3] Horwitz LI, Bradley EH., "Percentage of US emergency department patients seen within the recommended triage time: 1997 to 2006," *Arch Intern Med.*, November 9, 2009; 169 (20): 1857-1865.
- [4] Holroyd, B., Bullard, M., Latoszek, K. et al., "Impact of Triage Liaison Physician on Emergency Department Overcrowding and Throughput: A Randomized Controlled Trial," *Acad Emerg Med* 2007; Vol 14 (8): 702-708.
- [5] Russ, S., Jones, I, Aronsky, D, Dittus, RS, Slovis, C., "Placing Physician Orders at Triage: The effect on length of stay," *Annals of Emergency Medicine*, July 2010. Vol 56 (1): 27-33.
- [6] Tracy, M., "Triage Successes: A Hospital's Journey of Change and Growth," *Journal of Emergency Nursing*, Vol 33 (3), pp. 297-299. 2007.
- [7] Handel et al., "Emergency Department Throughput, Crowding, and Financial Outcomes for Hospitals," *Acad Emerg Med.*, August 2010; 17 (8): 840-847.
- [8] Fitzgerald, G., Jelinek, G., and Scott, D. et al., "Emergency department triage revisited," *Emerg Med J* 2010 27: pp 86-92, 2010.
- [9] Weiss et al., "Estimating the Degree of Emergency Department Overcrowding in Academic Medical Centers: Results of the National ED Overcrowding Study (NEDOCS)," *Academic Emergency Medicine*. Vol 11 (1): 39-50, January 2004.
- [10] Bernstein, S. et al., "The Effect of Emergency Department Crowding on Clinically Oriented Outcomes," *Academic Emergency Medicine*, Vol 16 (1): 1-10, January 2009.
- [11] Institute of Medicine. "Hospital-based Emergency Care: At the Breaking Point," Washington, DC: National Academies Press, 2006.
- [12] Bitterman, RA. "ED Triage—The New Hotbed of Litigation?" May 1 2009; ED Legal Letter Accessed August 27, 2010 www.allbusiness.com/health-care/health-care-facilities-hospitals/12580711-1.html
- [13] Lopez, L., et al. "Racial and Sex Differences in Emergency Department Triage Assessment and Testing Ordering for Chest Pain, 1997-2006." *Acad. Emer. Med.*, Vol 17 (8): 801-808, August 2010.
- [14] Gilboy N, Tanabe P, Travers DA, Rosenau AM, Eitel DR. *Emergency Severity Index, Version 4: Implementation Handbook*. AHRQ Publication No. 05-0046-2, May 2005. Agency for Healthcare Research and Quality, Rockville, MD. <http://www.ahrq.gov/research/esi/>
- [15] Aronsky, D. et al. "An Integrated Computerized Triage System in the Emergency Department." *AMIA 2008 Symposium Proceedings* 16-20.
- [16] Fitzpatrick, P., "Towards long-lived robot software," *Workshop on Humanoid Technologies, Humanoids 2006*, 2006.
- [17] Fitzpatrick, P., Metta, Natale, "Towards long-lived robot genes," *Robotics and Autonomous Systems*, 56(1):29-45, 2008.
- [18] Baars, B, & Franklin, S. "How conscious experience and working memory interact." *Trends in Cognitive Sciences* Vol. 7 No. 4, 2003.
- [19] K. Kawamura, R.T. Pack, M. Bishay and M. Iskarous, "Design philosophy for service robots" *Robotics and Autonomous Systems*, 18, 1996, 109-116.
- [20] Franklin, S., & Patterson, F. G. J. "The LIDA Architecture: Adding New Modes of Learning to an Intelligent, Autonomous, Software Agent," *IDPT-2006 Proceedings (Integrated Design and Process Technology)*: Society for Design and Process Science, 2006.
- [21] Zhang, Z., Dasgupta, D., & Franklin, S. "Metacognition in Software Agents using Classifier Systems," *Proceedings of the Fifteenth National Conference on Artificial Intelligence* (pp. 83–88). Madison, Wisconsin: MIT Press, 1998.
- [22] J. Piaget, *The Origins of Intelligence in Children*, (International University Press, 1952).
- [23] Gazzaniga, M.S. *Cognitive Neuroscience: The Biology of the Mind*, 2nd edn, (Norton and Co., NY, 2002).
- [24] Braver, T.S. and Cohen, J.D. "On the control of control: The role of dopamine in regulating prefrontal function and working memory," in *Control of Cognitive Processes: Attention and Performance XVIII*, eds., S. Monsell and J. Driver, (MIT Press, Cambridge, MA, 2000), pp. 713-738.
- [25] Miller, E.K. "Cognitive control: Understanding the brain's executive," *Fundamentals of the Brain and Mind*, Lecture 8, (MIT Press, Cambridge, MA, 2003).
- [26] Hommel, B. Ridderinkhof, K.R. and Theeuwes, J. "Cognitive control of attention and action: Issues and trends," *Psychological Research*, 66 (2002) 215-219.
- [27] Kawamura, K. and Gordon, S. "From Intelligent Control to Cognitive Control", 11th *International Symposium on Robotics and Applications (ISORA)*, July 2006
- [28] Kawamura, K. Gordon, S. Ratanaswasd, P. Erdemir E. and Hall, J. "Implementation of Cognitive Control for a Humanoid Robot", *International Journal of Humanoid Robotics*, vol. 5, no. 4 (2008), 547-586.
- [29] Endsley, M. R. "Toward a theory of situation awareness in dynamical systems." *Human Factors*, 37(1) 32-44, 1995.
- [30] Ramamurthy, U., D'Mello, Sidney, K., & Franklin, S. (2006, October 2006). "LIDA: A Computational Model of Global Workspace Theory and Developmental Learning." Paper presented at the *BICS 2006: Brain Inspired Cognitive Systems*.
- [31] Baars, B. J., & Franklin, S. "Consciousness is computational: The LIDA model of Global Workspace Theory." *International Journal of Machine Consciousness*, 1(1), 23-32, 2009.
- [32] Baars, B. *A Cognitive Theory of Consciousness*. Cambridge: Cambridge University Press, 1988.
- [33] Baars, Bernard J. "The conscious access hypothesis: origins and recent evidence." *Trends in Cognitive Science*, 6, 47–52, 2002.
- [34] Franklin, S., & Graesser, A., 1997. "Is it an Agent, or just a Program?: A Taxonomy for Autonomous Agents." *Proceedings of the Third International Workshop on Agent Theories, Architectures, and Languages, published as Intelligent Agents III*, Springer-Verlag, 1997, 21-35.
- [35] Franklin, S., Baars, B. J., Ramamurthy, U., & Ventura, M. "The Role of Consciousness in Memory." *Brains, Minds and Media*, 1, 1–38, 2005.

- [36] Wagner, A.R. and Arkin, R.C. "Analyzing Social Situations for Human-Robot Interaction." *Interaction Studies*, 10(2), 2008.
- [37] Franklin, S., Kelemen, A., & McCauley, L. "IDA: A Cognitive Agent Architecture," *IEEE Conf on Systems, Man and Cybernetics* (pp. 2646–2651): IEEE Press, 1998.
- [38] Franklin, S., & McCauley, L. "Interacting with IDA." In H. Hexmoor, C. Castelfranchi & R. Falcone (Eds.), *Agent Autonomy* (pp. 159–186). Dordrecht: Kluwer, 2003.
- [39] Garg, Amit X., Adhikari, Neill K. J., McDonald, Heather et al., "Effects of Computerized Clinical Decision Support Systems on Practitioner Performance and Patient Outcomes: A Systematic Review," *JAMA* (March 9, 2005), 293(10):1223-1238.
- [40] Franklin, S., & Jones, D. *A Triage Information Agent (TIA) based on the IDA Technology* Paper presented at the AAAI Fall Symposium on Dialogue Systems for Health Communication. Washington, DC, USA, 2004.
- [41] Hoiem, D., Efros, A. A., & Hebert, M. *Closing the Loop on Scene Interpretation*. Paper presented at the Conference on Computer Vision and Pattern Recognition (CVPR), 2008.
- [42] McCall, R., Franklin, S., Friedlander, D., & D'Mello, S. *Grounded Event-Based and Modal Representations for Objects, Relations, Beliefs, Etc.* Paper presented at the FLAIRS-23, Daytona Beach, FL. 2010.
- [43] Sulfaro, S. *Charting the Course for Triage Decisions*, *Journal of Emerg Nurs* 35:268-9, 2009.