

# An Architecture for Rehabilitation Task Practice in Socially Assistive Human-Robot Interaction

Ross Mead, Eric Wade, *Member, IEEE*, Pierre Johnson, Aaron St. Clair, Shuya Chen, and Maja J Matarić, *Senior Member, IEEE*

**Abstract**—New approaches to rehabilitation and health care have developed due to advances in technology and human robot interaction (HRI). Socially assistive robotics (SAR) is a subcategory of HRI that focuses on providing assistance through hands-off interactions. We have developed a SAR architecture that facilitates multiple task-oriented interactions between a user and a robot agent. The architecture accommodates a variety of inputs, tasks, and interaction modalities that are used to provide relevant, real-time feedback to the participant. We have implemented the architecture and validated its technological feasibility in a small pilot study in which a SAR agent led three post-stroke individuals through an exercise scenario. In the following, we present our architecture design, and the results of the feasibility study.

## I. INTRODUCTION AND BACKGROUND

THERE is increasing demand for innovative health-care interventions. The number of individuals requiring intense, personalized care is quickly outpacing the number of trained clinicians. A novel approach to remedying this problem comes from human-robot interaction (HRI). A subcategory of HRI applications focuses on using robots to provide interventions through hands-off, social interaction—this field is known as socially assistive robotics (SAR) [1]. The goal of SAR is to provide populations suffering from social and/or cognitive deficits with assistance through non-contact interaction, such as verbalization and gestures.

A group that stands to benefit significantly from SAR is the stroke population. In the US alone, over 800,000 people suffer from stroke each year. Of these, 400,000 go on to live with motor task deficits of some sort [2]. A large percentage of this population remains unable to voluntarily use the paretic arm, even though the ability to do so was identified as the most important outcome of self-reported recovery [3, 4, 5]. Task-oriented training (TOT) has been considered the primary approach to motor restoration for people post-stroke [6]. TOT is a treatment approach that addresses activity limitations by introducing practical daily activities to regain mobility and re-acquire skills [6]. The first nationwide, single-blind, randomized clinical trial of an intensive CIMT (a special class of TOT) training program demonstrated the beneficial effect on voluntary arm use for the patients with

hemiparesis [7], but it may not be realistic for an individual therapist to administer CIMT for the suggested six hours per day in the current clinical setting [6]. Socially assistive robots may help to fill this gap in care by enabling patients post-stroke to practice supervised TOT at clinics or at home.

Two crucial components of TOT are the responsiveness to the patient and the reliance on varied activity of daily living (ADL) tasks. Thus, our SAR architecture design requires: (1) the ability to accommodate varied ADL-inspired tasks without significant reconfiguration and (2) the provision of real-time task-dependent feedback to the patient/user. While these requirements are specific to our architecture, inspiration can be taken from related work. Robot Operating System (ROS; <http://www.ros.org/wiki/>) is an open-source framework for device intercommunication, utilizing distributed nodes interacting via a publish/subscribe paradigm to provide ease of use and portability. The Player project (<http://playerstage.sourceforge.net/>) shares this approach for device communication. Similarly, we have determined that a distributed architecture would be ideal for task-oriented training scenarios, considering the variability of inputs, tasks, and interaction modalities required for TOT. B<sup>3</sup>IA is a behavior-based SAR architecture designed for social skills intervention in children with autism; the hallmark of the architecture is its ability to monitor user activity over time to provide more meaningful interaction behaviors [8]. Likewise, we must manage constantly evolving user information during each interaction to better address the challenges of task-oriented training.

In the following, we describe our architecture and system in detail, present the results of a validation pilot study with post-stroke participants, and discuss continuing research.

## II. REHABILITATION TASK PRACTICE ARCHITECTURE

Our socially assistive human-robot interaction architecture guides a user through a series of rehabilitative tasks. It processes data from distributed user/world/robot activity servers and, upon recognizing a particular event or pattern within the context of the current task, requests the appropriate robot feedback action [Fig. 1].

### A. Task Manager

The *task manager* is the entry point and main system client of the rehabilitation task practice architecture. It provides the necessary interfaces (sockets) for all activity servers, and handles the context-switch between tasks, passing control to the appropriate task-oriented controller.

Manuscript received March 12, 2010. This work was supported in part by the National Science Foundation (NSF) Graduate Research Fellowship Program, as well as NSF Grants CNS-0709296 and IIS-0713697.

R. Mead, E. Wade, P. Johnson, A. St. Clair, S. Chen, and M. J. Matarić are with the University of Southern California, Los Angeles, CA 90089 USA (phone: 213-740-6245; fax: 213-821-5696; e-mail: {rossmead, ericwade, pierrejo, astclair, shuyache, mataric}@usc.edu).

By default, the manager contains robot behaviors for the introductory phase of an interaction between the robot and the human user, and utilizes logged user data to adapt the way the robot communicates with the user in subsequent interactions. Task metadata are also used to provide ties between tasks and produce context for task transitions.

In addition, the manager is responsible for maintaining a *world state*. This state is treated as common across all tasks.

### B. Task-Oriented Controllers

Each *task-oriented controller* guides the user through a single rehabilitation scenario, and is defined by metadata, user and robot states, and task-inspired conversational feedback.

*Task metadata* refer to communicative content that the manager uses to introduce or transition in-to/out-of the task. This information aims to provide the user with a common ground context for the scenario [9].

Each objective of a particular task is defined by a desired *task-oriented user state* in a finite state machine. State

transitions are governed by the appropriately specified user activities passed-in by the task manager. Similarly, *task-oriented robot states* represent collaborative human-robot objectives, as well as interaction monitoring and response modes of the robot.

The robot provides *task-inspired conversational feedback* to the user in response to task-oriented user activities and transitions. The communicative content is situated in the social context, allowing for “natural” interactions.

The specification of this information is designed to be simple and efficient. An extensible collection of task-oriented controllers has been developed (discussed in Section III-C). Controllers can be reordered within a program, allowing for dynamic interactions with the user.

### C. Activity Layers

The rehabilitation task practice architecture interacts with three layers of activity: (1) the *robot activity layer*, (2) the *user activity layer*, and (3) the *world activity layer* [Fig. 2]. Each layer consists of any number of *activity servers*, which

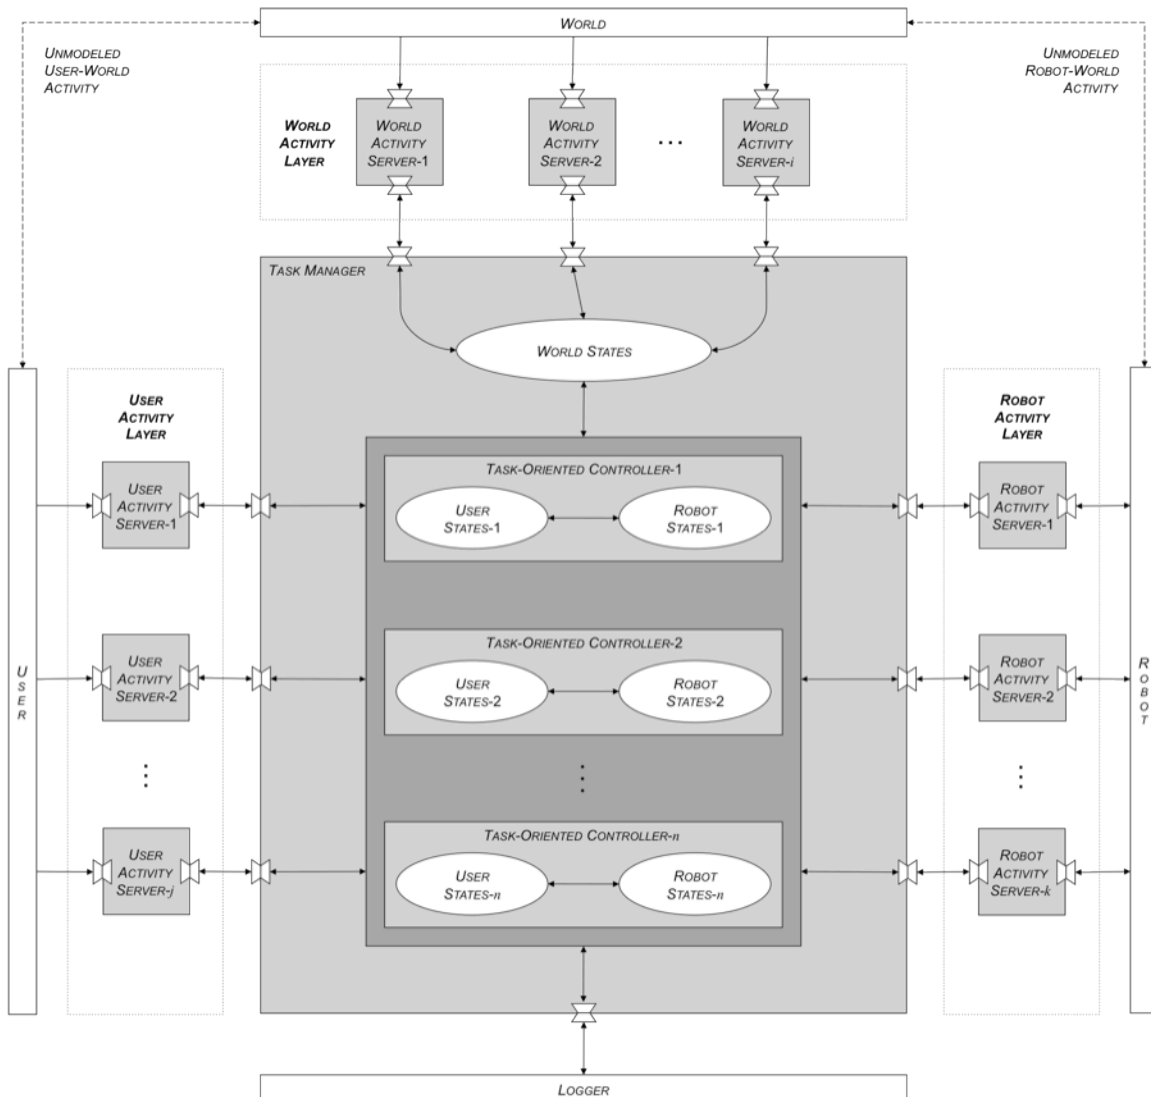


Fig. 1: The rehabilitation task practice architecture.

provide subsystem input (e.g., user gestures, location of an object, etc.) and/or output (e.g., robot verbal/nonverbal feedback).

A *robot activity server* (RAS) in the robot activity layer provides both request and response services. An RAS can communicate robot sensor information to a client process and/or can accept commands from the task manager for task-inspired conversational feedback from the robot.

Each *user activity server* (UAS) in the user activity layer communicates data from a user-wearable, environmental, or robot sensor to a client process. The task manager uses this information to model the state and progression of the user within the context of the current task. Additionally, a UAS can provide input to other UAS's for more complex user monitoring, such as nonverbal clusters, coverbal communication, and higher-level multimodal social interaction [10]. Changes in user activity often trigger transitions in task-oriented user state, thus, prompting task-inspired conversational feedback from the robot.

Processed information that does not directly pertain to the robot or the user is maintained by a *world activity server* (WAS) in the world activity layer. The task manager uses communicated world activities to maintain a consistent world state. This is particularly important in the context of post-stroke rehabilitation and TOT, as domain tasks often require the user to manipulate the environment. Monitored changes in world activities usually provide the most drastic events in an interaction, often signifying task completion and/or context shifts.

#### D. Logger

The *logger* maintains a history of all interactions with activity servers and task-oriented controllers. This information is used for performance analysis, as well as to influence subsequent communicative behaviors with the user. For example, a robot that has interacted with a specific user before need not reintroduce itself. Such user history is vital to developing persistent and meaningful relationships between humans and robots.

### III. IMPLEMENTATION

We have implemented the described architecture in a SAR framework for experimental analysis within the context of post-stroke rehabilitation. We used Player device interfaces (<http://playerstage.sourceforge.net/>) to interact with all hardware discussed below. We used Google's Protocol Buffers (<http://code.google.com/p/protobuf/>) for safe-socket communication between the task manager and all subsequent activity servers.

#### A. Robot Platform

The experimental test-bed used in this work was the Bandit III robot platform developed in USC's Interaction Lab, both physically and in simulation [Fig. 2]. Bandit is an upper-torso humanoid robot with 19 degrees of freedom: 7 in each arm (shoulder forward and backward, shoulder in and out, elbow tilt, elbow twist, wrist twist, wrist tilt,

grabber open and close; left and right arms), 2 in the head (pan and tilt), 2 in the lips (upper and lower), and 1 in the eyebrows. These high degrees of freedom allow the robot to be expressive by using individual and combined motions of the head, face, and arms. An extensive gesture and facial expression library has been developed to enhance the interactive experience. We are particularly interested in utilizing Bandit's anthropomorphic—but not highly realistic—appearance as a means of establishing user engagement, and comparing its impact to our prior work with non-anthropomorphic robot test-beds [11, 12, 13, 14]. The robot is more to scale with respect to human users than many other humanoid platforms; mounted atop a Pioneer P2 base, the entire robot is one meter tall, making it well suited for interaction. (For safety reasons, the robot was placed outside of the experiment workspace. In addition, a large 'kill-switch' button is built into the robot for use in case of any emergency.)



Fig. 2: The Bandit III upper-torso humanoid robot platform atop a Pioneer P2 base; physically embodied (left) and simulated (right).

#### B. Pilot Study

In the development of the post-stroke rehabilitation SAR system, we collaborated with physical therapists at USC's Health Sciences Campus (USC-HSC). The therapists evaluated the usability of the system; subsequently, based on their recommendations, we selected a series of tasks to evaluate the efficacy of the TOT architecture for motor-task rehabilitation with stroke patients. We were interested in comparing the relative efficacy and psychometric characteristics of a physically embodied robot to a 3D simulation [Fig. 2]. Other than their physical embodiments, the robot and simulation agents were identical. The tasks that the study participants performed, under the supervision of one of the agents, were: (1) lifting books from a desktop to a raised shelf, (2) moving pencils from one bin to another, and (3) opening and closing a jar.

#### C. Activity Layers

World activities (e.g., the placement of a book on a shelf) were detected through the use of a Nintendo Wiimote™ and a digital scale. Participants were outfitted with an inertia-based motion-capture (mocap) system, which monitored the person's arm and trunk (torso) movements. Each participant was led through assessment and rehabilitation activities by the robot using a combination of verbal and nonverbal interaction modalities. [Fig. 3]



Fig. 3: Bandit robot conducting a post-stroke rehabilitation session.

1) *Wiimote™ World Activity Server*: A Nintendo Wiimote™ was used to remotely start, pause, change, and stop an interaction, as well as monitor task-oriented button presses. Wiimote™ data is communicated via Bluetooth to a “Wiimote” world activity server (WM-WAS), which passed information to the task manager for input and logging.

2) *Object Transfer World Activity Server*: A United Nuclear UN-6000 digital scale (<http://unitednuclear.com/>) was used to detect when a participant interacted with the physical environment within the context of a particular task (e.g., the placement of a book on a shelf or pencil in a bin). This device provides a continuous data stream to an “object transfer” world action server (OT-WAS) over a USB-to-Serial connection. When an item (e.g., a book) is appropriately placed on the scale, the OT-WAS reports the change to the task manager, prompting task-inspired conversational feedback from the robot.

3) *Mocap Gesture User Activity Server*: Using inertial measurement units (IMUs), we developed a user-wearable motion capture system to monitor three-dimensional arm and trunk (torso) kinematic information [15]. The system consists of a 3-axis accelerometer, three single-axis rate gyros (one for each axis: x, y, and z), and a 3-axis magnetometer; it provides nine channels of 10-bit motion data over an I<sup>2</sup>C bus at 40 Hz per channel. The hardware is enclosed in a custom-molded, lightweight plastic casing attached to an adjustable Velcro armband. User motion data were windowed and processed “on-board” by a small user-wearable computer system—a Gumstix-Robostix-Wifi pack, a battery, and power regulation components (<http://www.gumstix.com/>). Processed motion of the user is converted into one of a series of learned task-relative gesture symbols (e.g., “arm lifting”, “arm lowering”, “arm moving too slow”, “trunk leaning forward”, “trunk twisting”, etc.) and communicated wirelessly to the robot via the task manager. The entire mocap system is small, lightweight, and inexpensive, and is considered a non-invasive and non-significant-risk device.

A “mocap gesture” user activity server (MG-UAS) was developed that provides task-related gesture symbols [15]. When the task manager receives a gesture symbol (from the

MG-UAS) that is valid within the context of the current task, it generates task-inspired conversational feedback (scripted in this implementation), attempting to motivate and/or improve user performance. It is, however, easy to “cheat” the mocap system alone (e.g., by moving the arm without lifting a book) and, thus, these symbols are fused with data from the OT-WAS to produce appropriate robot actions.

4) *Verbal Robot Activity Server*: User and world activities communicated by the servers often trigger requests (by the task manager) for robot actions from “verbal” and “nonverbal” robot activity servers (V-RAS and NV-RAS, respectively) in the form of task-inspired conversational feedback. For this implementation, a collection of over 400 phrases were carefully scripted (based on the task, user status, robot personality, etc.) and recorded (by a human voice actor). Each phrase was considered to have a valence—for our purposes, an intrinsic positive, negative, or neutral attitude/meaning—associated with it. For example, a positive valence phrase we used was “Great job! Continue to work this hard and you’ll recover faster.” Conversely, a negative valence phrase we used was “Do not throw the book onto the shelf.” Phrases with neutral or no valence were also used, such as “The new exercise consists of lifting a book from a desktop to a raised shelf.” The method for extracting phrase valence is discussed below.

5) *Nonverbal Robot Activity Server*: The controller can request one of two types of nonverbal robot actions from the NV-RAS: speech-independent gestures (illustrators) and speech-dependent (coverbal) behaviors [10]. The first of these nonverbal actions—illustrators—were hard-coded, and were used to provide task instruction or motivation to the user by demonstration from the robot. For example, if the robot detects (through feedback from the MG-UAS) that the user is not moving his/her arm sufficiently for the task, it acts out the task itself at the proper speed, thus, attempting to encourage the user to match the rate of movement activity of the robot.

The second of these nonverbal actions—coverbal behaviors—attempt to synchronize robot gestures to the verbal output (i.e., the NV-RAS was also a client of the V-RAS). To determine which gestures correspond to verbal content, we parsed the scripts that were used for the audio recordings using the NonVerbal Behavior Generator (NVBG) rules described in [16]. These rules are based on specific keywords/structures and behavior priorities learned from user activity modeling. This approach is simple to implement; however, its reliance on individual words or structure in the phrase means that the each word in the phrase must be time-annotated. In our initial implementation, this information was not available. As an alternative, we determined the duration of the entire length of a phrase (stored as an audio file) and generate an overall “phrase valence”, which each dictated the duration and type of nonverbal behavior exhibited by the robot for that particular phrase utterance, respectively.

The overall valence of a phrase is determined based on all of the NVBG rules that have been applied to that phrase. For the purposes of this study, a single NVBG rule,  $R$ , contains a

label,  $L$ ; a list of keywords,  $K$ ; a priority (from [16]; lower value = higher priority),  $P$ ; and a valence,  $V$ . The valence parameter augments the standard definition of an NVBG rule described in [16], providing a sense of emotion in the nonverbal display. The suggested output behavior of each rule affects a certain number of words,  $n$  ( $1 \leq n \leq N$ , where  $N$  is the number of words in the phrase); in most cases, the nonverbal behavior is demonstrated either throughout the entirety of the phrase (i.e.,  $n = N$ ) or co-occurring with the keyword(s) associated with the rule (i.e.,  $n = 1$  or  $2 \dots$ ). If  $m$  NVBG rules are applied to a phrase, then the weighted valence,  $v_i$  (based on the proportion of words affected), of the  $i^{\text{th}}$  rule ( $1 \leq i \leq m$ ) applied to the phrase can be written as:

$$v_i = \frac{n_i V_i}{N P_i} \quad (1)$$

Thus, the weighted valence of each NVBG rule applied to the phrase can then be accumulated to determine the overall phrase valence,  $V^*$ :

$$V^* = \frac{\sum_{i=1}^m v_i}{\sum_{i=1}^m \frac{n_i}{N}} = \frac{\sum_{i=1}^m \frac{n_i V_i}{N P_i}}{\sum_{i=1}^m \frac{n_i}{N}} = \frac{\sum_{i=1}^m \frac{n_i V_i}{P_i}}{\sum_{i=1}^m n_i} \quad (2)$$

For phrases with positive valence ( $V > 0$ ), the robot nods its head with a valence-proportional amplitude at regular intervals throughout the entire phrase. For phrases with negative valence ( $V < 0$ ), the robot shakes its head with a valence-proportional amplitude at regular intervals throughout the entire phrase. There is no head movement for phrases with neutral valence ( $V = 0$ ). For all phrases, the hands are shown palms-up with a “beat” gesture [16]. Such nonverbal/affect behaviors will be extended (primarily based on time) in our continuing and future work.

#### IV. SUMMARY AND DISCUSSION

In the pilot study, three post-stroke individuals in the chronic stage of recovery practiced the book-shelving task using our SAR architecture. The upper extremity functional motor ability of the participants was characterized using the Fugl-Meyer Assessment (FMA). The participants had FMA scores of 48, 45, and 46. They came to the USC-HSC Health Sciences Campus for a 2-session experiment. The first session consisted of assessments administered by our team clinicians, and the second consisted of the task practice. During the task practice, the participants interacted with and received feedback from one of the two different SAR agents—a physically embodied robot and a simulated robot [Fig. 2], all participants experienced both agents. In both cases, the SAR agent monitored the participant through the sensors described in Section III-C, and used those inputs to verbalize and gesture in response to user performance.

#### A. Survey Results

At the end of the final session, we administered a 12-item, likeability/preference questionnaire that elicited user feedback regarding the quality, quantity, and content of the SAR system behavior. The survey also included questions regarding the participant’s previous exposure to technology, preference of agent embodiment, and perception of the agents’ gestures. The analysis of the questionnaire data revealed that the participants preferred the physical robot. Of the six questions regarding preference, 2 participants rated the physical robot more highly for 5 questions, and had no preference on the sixth. The third participant had no preference. Because the task was identical with both SAR agents, we suspect the embodiment was the cause of the preference for the physical robot. However, with the small sample size, we cannot be conclusive about what characteristic of the embodiment was preferred by the participants. In Section V we discuss how the questionnaire results provide insights for subsequent studies.

#### B. Robustness and Fault-Tolerance

One important outcome of our pilot study was an empirical determination of the robustness of our architecture. Due to the multiplicity of its interacting subsystems, it is critical that unforeseen occurrences in one part of the system do not cause the entire system to shut down. We were able to observe the architecture’s ability to continue functioning despite failure on three occasions.

In one instance, the motion suit cable was pulled out by a participant’s vigorous motion, causing the motion suit’s I<sup>2</sup>C communication bus to drop. This resulted in the controller receiving symbols that indicated the participant was no longer moving. However, because the scale continued to provide accurate data (registering the occasional shelving of magazines), the SAR agent continued to provide appropriate feedback, and the interaction continued without the participant noticing any drastic changes.

In another instance, the serial cable that connected the robot Bandit to the nonverbal robot activity server was not adequately secured and the robot lost its ability to perform physical gestures. However, it maintained its ability to provide verbal feedback to the participant. The activity continued for the allotted session time. Afterwards, the participant did notice that the robot “seemed stiff,” but otherwise thought the interaction was normal.

Finally, in another instance, the scale shut down unexpectedly during an experiment. In this case, the architecture could no longer recognize the magazines being shelved, but did recognize that the user was moving. This allowed for the continued provision of feedback regarding motion (but less congratulatory feedback).

In each of these cases, the interaction session continued as the architecture continued to function. While some redundancy of sensing is required (in this case, both the scale and motion suit being used to monitor performance), the key observation based on these errors is that the

architecture is capable of maintaining the interaction in spite of individual subsystem failures.

## V. ONGOING AND FUTURE WORK

In addition to demonstrating the technical feasibility of our task practice architecture, the pilot study also provided a number of valuable insights that have led us to refine our experimental design. This has necessitated further expansion in the number and type of sensors employed in the current framework for monitoring user and world activity.

The ongoing follow-up experiment examines the importance of maintaining challenge level in post-stroke rehabilitation by using a custom-built wire puzzle (a piece of heavy-gauge wire and assorted ring-shaped wands). Subjects are instructed to pass the ring along the wire without making contact between the two. The wand and the wire puzzle are connected to the digital inputs of a USB Phidget Interface Kit allowing for the system to monitor user progress. The Phidget Kit controller (<http://www.phidgets.com/>) functions as a data acquisition board. In addition, we have added sensors for more precise monitoring of the participant. The current mocap system is combined with a 6-IMU data glove (<http://www.aceleglove.com/>) that will allow analysis of finger(s), hand, and arm motion; a Wiimote™ with MotionPlus™ attachment will be used to derive precise wand movements. These sensors are being incorporated into user activity servers that will derive the kinematics and dynamics of the affected limb in rehabilitation exercises and allow for appropriate corrective action by the robot and fine-grained data logging.

To facilitate tasks that require manipulation of multiple objects, we have implemented an overhead, camera-based tracking system as an “object pose” world activity server (OP-WAS). To simplify the segmentation and tracking problem, tracked objects are augmented with one or more fiducial markers, which are then tracked using ARToolKitPlus [17] to accurately identify and detect the poses of each marker in the scene. The OP-WAS then supplies a unique identification number and pose information, consisting of a 3D position and yaw/pitch/roll for each detected tag. The presence or absence of a tag at a particular instant in time can be used by the system to make inferences regarding the occlusion characteristics of the current object configuration (e.g., has an object been stacked atop another object). A similar system is also employed as a “head pose” user activity server (HP-UAS), where a hat tagged with a marker is used to estimate the head position and orientation of the user. This can be combined with object tracking to estimate the scope of visual attention from the user's current point of view.

The robot dialogue has also been streamlined via the integration of a commercial text-to-speech engine (<http://www.neospeech.com/>). This simplifies subsequent implementations of verbal utterances and eliminates the dependence on prerecorded human speech, while keeping the flexibility of the current verbal activity system.

## ACKNOWLEDGMENTS

The authors would like to thank Dr. Adriana Tapus for her preliminary work on the architecture, Cynthia Kushi for her help in administering the pilot study, and Jina Lee and Dr. Stacy Marsella for their cooperation in working with NVBG.

## REFERENCES

- [1] D. Feil-Seifer and M. J. Matarić, “Defining socially assistive robotics,” *Proceedings of the International Conference on Rehabilitation Robotics (ICORR-05)*, Chicago, Illinois, 2005.
- [2] AHA Statistical Update, *Circulation*. 2009; 119:e21-e181.
- [3] R. N. Barker R.N. and S.G. Brauer, “Upper limb recovery after stroke: the stroke survivors perspective,” *Disability & Rehabilitation*. Vol. 27, 2005, pp. 1213-1223.
- [4] R. N. Barker, T. J. Gill, and S. G. Brauer. “Factors contributing to upper limb recovery after stroke: a survey of stroke survivors.” *Disability & Rehabilitation*. Vol. 29, 2007, pp. 981-989.
- [5] N. E. Mayo, S. Wood-Dauphinee, R. Côté, L. Durcan, and J. Carlton. “Activity, participation, and quality of life 6 months poststroke.” *Archives of Physical Medicine and Rehabilitation* - August 2002 (Vol. 83, Issue 8, Pages 1035-1042, DOI: 10.1053/apmr.2002.33984).
- [6] N. Schweighofer, C. E. Han, S. L. Wolf, M. A. Arbib, and C. J. Winstein, “A functional threshold for long-term use of hand and arm function can be determined: predictions from a computational model and supporting data from the Extremity Constraint-Induced Therapy Evaluation (EXCITE) Trial.” *Phys Ther*. 2009 Dec;89(12):1327-36. Epub, Oct 2009.
- [7] S. L. Wolf, C. J. Winstein, J. P. Miller, E. Taub, G. Usatte, D. Morris, C. Giuliani, K. E. Light, and D. Nichols-Larsen. “Effect of constraint-induced movement therapy on Upper Extremity Function 3 to 9 Months After Stroke: The EXCITE Randomized Clinical Trial.” *JAMA*. 2006;296:2095-2104.
- [8] D. J. Feil-Seifer and M. J. Matarić, “B31A: an architecture for autonomous robot-assisted behavior intervention for children with autism spectrum disorders,” *Proceedings of the IEEE International Workshop on Robot and Human Interactive Communication*, Munich, Germany, Aug 2008.
- [9] S. Kiesler, “Fostering common ground in human-robot interaction,” *Proceedings of the 14th IEEE International Workshop in Robots and Human Interactive Communication (RO-MAN '05)*, pp. 729-734, IEEE Press, 2005.
- [10] M. L. Knapp and J. A. Hall, *Nonverbal Communication in Human Interaction*, 7<sup>th</sup> edition, Boston, Massachusetts: Wadsworth Publishing, 2009.
- [11] J. Eriksson, M. J. Matarić, and C. Winstein, “Hands-off assistive robotics for post-stroke arm rehabilitation,” *Proceedings of the IEEE International Conference on Rehabilitation Robotics (ICORR-05)*, Chicago, Illinois, 2005.
- [12] R. Gockley and M. J. Matarić, “Encouraging physical therapy compliance with a hands-off mobile robot,” *Proceedings of the First International Conference on Human Robot Interaction (HRI-06)*, Salt Lake City, Utah, 2006.
- [13] A. Tapus and M. J. Matarić, “Towards socially assistive robotics,” *International Journal of the Robotics Society of Japan (JRSJ)*, vol. 24, no. 5, 2006, pp. 576-578.
- [14] A. Tapus and M. J. Matarić, “User personality matching with hands-off robot for post-stroke rehabilitation therapy,” *Proceedings of the 10th International Symposium on Experimental Robotics (ISER)*, Rio de Janeiro, Brazil, 2006.
- [15] E. Wade and M. J. Matarić. “Design and testing of lightweight inexpensive motion-capture devices with application to clinical gait analysis”. In *Proceedings of the International Conference on Pervasive Computing*, pages 1 - 7, Aug 2009.
- [16] J. Lee and S. Marsella, “Nonverbal behavior generator for embodied conversational agents,” *6<sup>th</sup> International Conference on Intelligent Virtual Agents*, Marina del Rey, California, 2006.
- [17] D. Wagner and D. Schmalstieg, “ARToolKitPlus for pose tracking on mobile devices,” *Proceedings of 12th Computer Vision Winter Workshop (CVWW'07)*, 2007, pp. 139-146.