

Robot Exercise Instructor: A Socially Assistive Robot System to Monitor and Encourage Physical Exercise for the Elderly

Juan Fasola and Maja J. Matarić, *Senior Member, IEEE*

Abstract— We describe the design and implementation of a socially assistive robot that monitors the performance of a user during a seated arm exercise scenario, with the purpose of providing motivation to the user to complete the task and to improve performance. The visual arm pose recognition procedure used by the robot in tracking user performance, the three exercise games, and the methodology behind the human-robot interaction dialogue are presented. A two-condition experimental study was conducted with elderly participants to test the feasibility and effectiveness of the robot exercise system, the results of which demonstrate the viability and usefulness of the system in motivating exercise among elderly users.

I. INTRODUCTION

SOcially Assistive Robotics (SAR) [1] is an area of Human-Robot Interaction (HRI) that focuses on aiding through social rather than physical interaction between the robot and the human user, and has the potential to enhance the quality of life for large user populations, including the elderly, people with physical impairments and those involved in rehabilitation therapy. The world's population is growing older, thereby introducing a wide array of societal challenges. It is estimated that in 2050 there will be three times more people over the age 85 than there are today. Many are expected to need physical and cognitive assistance. The demand for such assistance is quickly outpacing the available supply of available and affordable human care. As the elderly population continues to grow, a great deal of attention and research will be dedicated to assistive systems that allow the elderly to live independently in their own homes. As such, the main purpose of socially assistive robotics technology is not to replace human care-givers, but rather to provide assistance where human assistance is not available or affordable.

The robotics literature that addresses the area of assistive robots for the elderly is limited. Researchers have developed robots that focus on providing assistance for functional needs, such as mobility and health monitoring [4], navigation and schedule reminders [5], as well as social and emotional needs, such as reducing depression [6] and increasing social interaction [7].

Manuscript received March 12, 2010. Support for this research was provided by a grant from the Robert Wood Johnson Foundation's Health Games Research program and the NSF CRI CNS-0709296 grant.

J. Fasola is with the University of Southern California, Los Angeles, CA 90089 USA (e-mail: fasola@usc.edu).

M. J. Matarić is with the University of Southern California, Los Angeles, CA 90089 USA (e-mail: mataric@usc.edu).

Our focus is on designing and studying the effectiveness of socially assistive robotics technology towards providing affordable customized care for individuals in need of assistance [2], [3], with particular emphasis on how the user's intrinsic motivation can be influenced by a socially assistive robot in order to maximize the probability of success of the therapeutic intervention.

II. GOALS AND APPROACH

Regular physical exercise has been shown to be effective at maintaining and improving the overall health of elderly individuals [8], [9], [10], [11]. The work presented in this paper is an experimental implementation of a socially assistive robot whose purpose is to motivate users to exercise by engaging in a simple seated arm exercise scenario. The overall goal of the pilot study is to evaluate and validate the effectiveness of the approach in order to gain insight for future, long-term studies involving the intended user populations in need of such customized assistance, such as the elderly and/or stroke patients. The specific aims of the work were as follows:

- 1) To create a socially assistive robot capable of interacting with elderly users in an exercise scenario, while providing feedback, praise, and motivation throughout.
- 2) To conduct a study with elderly participants to test the feasibility and effectiveness of our system design, as well as to gain insight into the role of user autonomy and choice within the exercise scenario.
- 3) To get feedback from the participants for the purpose of improving the system design.

III. ROBOT EXERCISE SYSTEM

A. System Overview

The exercise scenario consists of a socially assistive robot whose purpose is to instruct, evaluate, and encourage users to perform simple arm gesture exercises. The scenario is one-on-one, allowing the robot to focus its attention on the single user in order to provide timely, accurate feedback, and to maximize the effectiveness of the exercise session for the user. In the set up, the user is seated in a chair in front of the robot; the user and robot face each other. A black curtain is used as a backdrop to facilitate the visual perception of the user's arm movements. The complete exercise setup is shown in Fig. 1(a).

During the exercise sessions, the robot asks the user to perform simple seated arm gesture exercises. The range of the robot’s arm motion in the exercises is restricted to the sides of the body in order to improve the accuracy of the robot’s visual detection of the user’s arms when repeating the user-demonstrated motions. This type of seated exercise, called “chair exercise” or “chair aerobics”, is commonly practiced in senior living facilities and provides grounding for our exercise system. Chair exercises are highly regarded for their accessibility to those with low mobility [8], [9], [10], [11], for their safety as they reduce the possibility of injury due to falling from improper balance [8], [11], and for their health benefits such as improved flexibility [8], [10], muscle strength [8], [10], [11], ability to perform everyday tasks [8], [10], [11], and even memory recall [9].

The user is able to communicate with the robot through the use of a wireless button control interface, the popular Wiimote remote control, which communicates via Bluetooth with the button labels modified to suit our exercise system. There are two buttons available to the user to respond to prompts from the robot, labeled “Yes” and “No”, and one button for the user to request a rest break at any time during the interaction.

B. Robot Platform

To address the role of the robot’s physical embodiment, we use a biomimetic anthropomorphic robot platform that consists of a humanoid torso mounted on a MobileRobots Pioneer 2DX mobile base. The torso comprises 19 controllable degrees of freedom, which include: 6 DOF arms (x2), 1 DOF gripping hands (x2), 2 DOF pan/tilt neck, 1 DOF expressive eyebrows, and a 2 DOF expressive mouth. A photograph of the robot is shown in Fig. 1(b).

A standard USB camera is located at the waist of the robot to capture the user’s arm movements, allowing the robot to provide appropriate feedback about the exercise.

The robot’s speech is generated by the commercially available NeoSpeech text-to-speech engine [19], and a speaker on the robot outputs the synthesized voice to the user. The robot’s lip movements are synchronized with the spoken words for increased naturalness of speech.

C. Exercise Games

Three exercise games are available in our system: the Workout game, the Imitation game, and the Memory game. During an exercise session, the user is given the opportunity to play all three games, and often any game more than once. The following is a description of each game in detail.

1) *Workout Game*: In this game, the robot fills the role of a traditional exercise instructor by showing the user which arm gesture exercises to perform, by demonstrating with its own arms, and asking the user to imitate. The robot gives the user feedback in real time, providing corrections when appropriate (e.g., “Raise your left arm and lower your right arm” or “Bend your left forearm inwards a little”), and praise in response to each successful imitation (e.g., “Great job!” or




Fig. 1: (a) Exercise setup with user and robot facing each other. (b) The socially assistive robot platform

“Now you’ve got the hang of it”). In monitoring user performance, the robot compares the user’s current arm angles as detected by the vision module to those of the specified goal arm angles to determine performance accuracy. The comparison procedure is robust to user fatigue and variations in range of motion; it relies more on the user’s current hand positions and forearm angles than on the absolute differences of the goal and user arm angles.

2) *Imitation Game*: In this game, the roles of the user and robot from the Workout game are reversed, and thus the user becomes the exercise instructor showing the robot what to do. The robot encourages the user to create his/her own arm gesture exercises, and imitates user movements in real time.

As the roles of the interaction are reversed, with the robot relinquishing control of the exercise routine to user, the robot no longer provides instructive feedback on the exercises. However, the robot does continue to speak and engage the user by means of encouragement and general commentary. For example, if the robot detects that the user is not moving her arms, it encourages the user again to create new gestures by saying, for instance, “Mary, try and come up with your own gestures and I’ll imitate you.” In addition, the robot makes general comments about the game or user, such as “You’re a good instructor Mary” or “This is my favorite game, thanks for the workout.”

3) *Memory Game*: In this game, the user is challenged to learn a sequence of different arm gestures. The goal of the game is for the user to try and memorize ever-longer sequences, and thus compete against his/her own high score. The sequence is determined at the start of the game and does not change for the duration of the game. The arm gesture poses used for each position in the sequence are chosen at random at run time, and the sequence is of infinite length, thereby making the game challenging for users at any skill level.

The robot starts out by showing the first two positions of the sequence and asks the user to perform them while it provides feedback. Once the user has successfully repeated the first two gestures with the help of the robot, the user is asked to repeat the sequence again from the beginning, this time without demonstration or verbal feedback from the robot. Once the gestures are completed without help, the

robot shows the next two gestures in the sequence, and the user is again asked to perform the entire sequence from the beginning (now four gestures in length). As the user continues to successfully memorize all shown gestures, the robot continues to show the user two more (next six, and then eight gestures in total length), and the game progresses in difficulty.

The robot helps the user to keep track of the sequence by counting along with each correct gesture, and reminding the user of the poses when it detects errors (e.g., “Oh, that’s too bad! Here again is gesture five”). The robot also reports to the user his/her current high score (i.e., the number of gestures remembered correctly) in an attempt to motivate improvement upon past performance.

D. General Interaction Methodology

Many social intricacies contribute to the foundation of a meaningful relationship, details of which Bickmore and Picard have outlined in regards to human-computer interaction [18] but which also apply to human-robot interaction. These include empathy, humor, reference to mutual knowledge, continuity behaviors, politeness, and trust, among others. We place great importance on these relationship building tools and integrated each, in some form or another, into the social interaction component of our robot exercise instructor.

Our primary focus was in eliminating the perceived repetitiveness of the robot’s verbal instructions/comments. We believe that if the robot is perceived by the user as repetitive and hence predictable, this can lead to a decrease in the perception of the robot’s intelligence by the user, and ultimately to a loss of trust in the robot’s helpfulness in motivating exercise. Special attention was therefore placed on adding variety to the robot’s utterances.

Towards this end, the robot always draws from a list of phrases that emphasize the same point when speaking to the user, choosing one randomly at run time. For example, there are more than ten different ways in which the robot can praise the user (e.g., “Awesome!”, “Nice job!”, “Fantastic!”). Furthermore, if the robot needs to repeat itself exactly, for example when providing the same feedback comment during the Workout game, it adds filler words to the given phrase, such as the user’s name or the word “try” or both (e.g., “Try and raise your left arm”, “John, raise your left arm”).

Adding the user’s name to the interaction dialogue was an important part of our system design, not only to add variability, but also for its relationship building abilities [18]. The robot always uses the user’s name at the first greeting, and also when bidding farewell at the end of a session.

IV. VISION MODULE

In order to monitor user performance and provide accurate feedback during the exercise routines, the robot must be able to recognize the user’s arm gestures. To accomplish this, we developed a vision module that recognizes the user’s arm




Fig. 2: (a),(c),(d) Example face and arm angle detection results superimposed over original grayscale camera frames. (b) Segmented image of camera frame shown in (a).

gestures/poses in real time, with minimal requirements for the surrounding environment and none for the user.

Several different approaches have been developed to accomplish tracking of human motions, both in 2D and 3D, including skeletonization methods [14], [15], gesture recognition using probabilistic methods [16], and color-based tracking [17], among others. We opted to create an arm pose recognition system that takes advantage of our simplified exercise setup in order to achieve real-time results without imposing any markers on the user.

To simplify the visual recognition of the user, a black curtain was used to provide a static and contrasting background for fast segmentation of the user’s head and hands, the most important features of the arm pose recognition task, independent of the user’s skin tone.

The arm pose recognition algorithm consists of the following four major steps:

- 1) *Create segmented image*: The original gray-scale camera frame is converted into a black/white image by applying a single threshold over the image. All pixels below the threshold are set to black, and the rest to white. Fig. 2(a) shows an original grayscale image captured from the camera, with the segmented image shown in Fig. 2(b).

- 2) *Detect the user’s face*: The OpenCV frontal face detector is used to determine the location and size of the user’s face. With these values, an estimate for the shoulder position for both sides of the body is made.

- 3) *Determine hand locations*: The hand locations of the user are determined by examining the extrema points of the body pixels (max/min white pixel locations along x and y directions away from body) inside the region above the chest line and to the side of the face in the segmented image. The algorithm applies a simple set of rules, or heuristics, to choose which extrema points correspond to the hand location

for the given arm. For example, if highest white pixel (body pixel) in the segmented image on the left side is further than the approximate shoulder-to-forearm distance, then that is the left hand location.

4) *Determine arm angles:* Once the hand location for a given arm is found, the elbow point is estimated, which in turn provides the desired arm angles. The elbow point is estimated using the white pixel (body pixel) that lies furthest from the line connecting the hand position and the shoulder position, while also not exceeding the maximum allowable distance from the hand (to enforce forearm length restriction). Example arm angle detection results can be seen in Fig. 2(a), (c) and (d).

The vision module only searches for frontal faces, and thus detection rates depend largely on whether or not the user is facing the robot. However, high detection rates are not actually necessary for accurate gesture recognition, as the most recent detected face location is used in the arm detection procedure if a face is not found in the current frame. This substitution works well, since the user's head position generally remains stationary while the user is seated in a chair in front of the robot throughout the interaction.

The arm pose recognition algorithm runs with an average frame rate of 20 fps on a 2.4 GHz Intel Core 2 Duo processor, and while the arm detection procedure was not formally evaluated, the system was observed to be fairly robust to different types of clothing, lighting, and user body types, with a notable arm angle estimation accuracy that is sufficient for the purposes of our exercise system.

V. STUDY

A. Method

A feasibility study was conducted with residents from our partner senior living facility, Southern California Presbyterian Homes. Each participant took part in an exercise session with the robot twice per week for two weeks, yielding a total of four exercise sessions per participant. The primary purpose of the study was to test and validate the effectiveness of our robot exercise instructor system as well as to obtain valuable feedback from the participants for possible system improvement. The study was also designed to explore the role of choice and user autonomy within the exercise game scenario. Fig. 3 shows a participant interacting with the robot during the study.

Self-determination, represented in a task such as choice of activity, has been shown to increase or be less detrimental to intrinsic motivation when compared to similar task conditions that do not involve choice [12], [13]. These results are interesting because increasing the user's intrinsic motivation to complete the task, or in our case to exercise, is perhaps the most effective way to encourage long-term behavior change.

B. Experiment Conditions

Two experiment conditions were created to test user preferences regarding choice of activity: the Choice



Fig. 3: Study participant creating a novel gesture for the robot to perform during the Imitation game.

condition, and the No Choice condition. The conditions differ only in the manner in which the exercise games (Workout, Imitation, Memory) are chosen to be played during the exercise sessions. The design of the experiment is within-subject, as each participant engages in both conditions at least twice. The following are descriptions of each condition:

1) *Choice Condition:* The user is given the choice of which game to play at specific points in the interaction. The robot prompts the user to press the "Yes" button when he hears the game he would like to play, as the robot calls out the names of each of the three game choices. After the user has made a choice, the chosen game is played for a duration ranging from one to two minutes in length, after which the robot asks the user if he would like to play a different game. Depending on the user's response, the robot either continues playing the same game for another one to two minutes, or prompts the user again to choose which game to play next.

2) *No Choice Condition:* The robot chooses which of the three games to play at the specified game change intervals (every one to two minutes). The robot always changes from one game to another to try to minimize the user's frustration, as the robot is unaware of the user's personal game preferences. For simplicity, in this condition the robot always chooses to first play the Workout game, followed by the Imitation and Memory games, then cycles through them in the same order again.

C. Hypotheses

The hypotheses for the study were as follows:

- 1) After the first session, no significant difference will be found between the participants' preferences for the Choice or No Choice conditions;
- 2) After the fourth and final session, more participants will prefer the Choice condition over the No Choice condition.

The first hypothesis is based on the notion that the participants, having just been introduced to the exercise system, will not yet have any strong feelings about any of the three available games, and therefore will have little to no preference to whether they themselves or the robot chooses which games to play.

The second hypothesis stems from the idea that the participants, having already interacted with the robot in exercise sessions for two weeks and having gained more experience with each game, will have developed personal game preferences and therefore will prefer to choose for themselves which games to play.

VI. RESULTS

A. Participant Statistics

Twelve participants were recruited to participate in the study, one of whom dropped out after two sessions due to scheduling conflicts. Thus, eleven participants were able to complete the entire four-session, two-week trial. There were ten female participants and one male participant. All of the participants were seniors over the age of sixty-five.

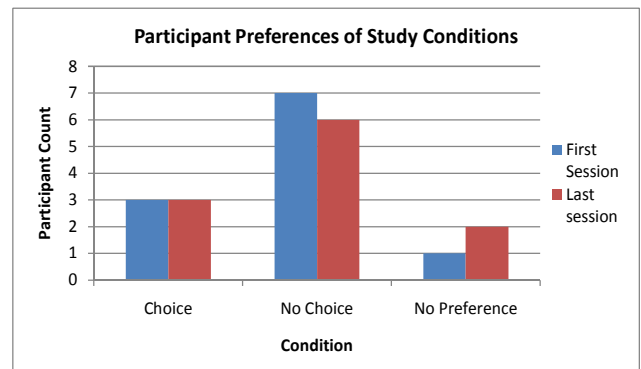
B. Participant Preferences Regarding Choice

The first exercise session presented each participant with both study conditions, each lasting ten minutes in duration. The order in which the conditions were presented was counter-balanced among the participants. For the next two sessions, the participants interacted with the robot only in the Choice condition, and for the fourth and final session, they interacted only in the No Choice condition.

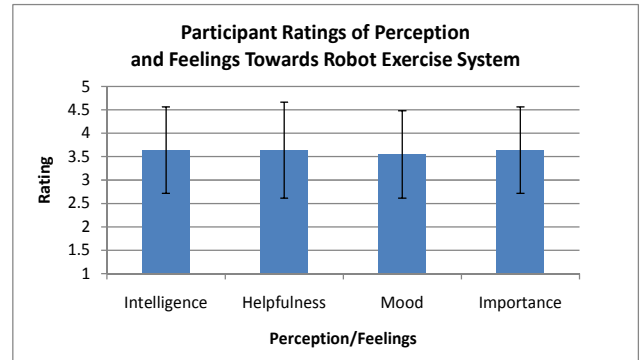
The survey results after the first session indicate that 3 participants preferred the Choice condition, 7 participants preferred the No Choice condition, and 1 participant had no preference. These results are consistent with our first hypothesis, as the difference in participant preferences between the two conditions is not statistically significant. It is interesting to note, however, that even though most participants preferred the No Choice condition, almost all of them, 10 out of 11, reported increased enjoyment of the task when given the opportunity to choose which exercise game to play; this result is consistent with the literature on choice effects on intrinsic motivation.

After the fourth and final session, the participants were again asked to state their preference among the two conditions, and the results are similar to those of the first survey. Two participants switched their preference from one condition to the other, and one additional participant expressed no preference. This result does not support our second hypothesis, and may be due to the fact that the enjoyment derived from choosing the game does not outweigh the enjoyment derived from relaxation due to reduced responsibility, as some of the participants reported being the case. A graph summarizing the participants' preferences of study conditions is shown in Fig.4 (a).

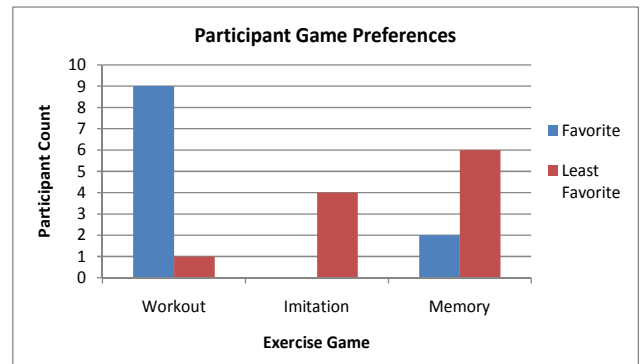
It is interesting to note that even though a majority of the participants reported preferring the No Choice condition after the final session, all 11 participants at one point or another during the study took advantage of having a choice in the Choice condition. Specifically, when given the option by the robot to change games, the participants either chose to continue playing the same game they were playing or chose



(a)



(b)



(c)

Fig. 4: (a) Graph of the participants' preferences of study condition, when surveyed after their first and last exercise session. (b) Graph of the participants' ratings in response to survey questions on their perception of the robot's intelligence, helpfulness, their mood during sessions, and how important the sessions were to them. (c) Graph of participants' preferences of exercise game.

to avoid playing a game they did not want to play. Neither of these cases would occur in the No Choice condition, as the robot was unaware of the user's current game preferences. This result speaks to the usefulness of user choice in the game scenario, suggesting that a hybrid approach that includes both user and robot decision making, tuned automatically for each user, might ultimately be best for achieving a fluid and enjoyable game interaction for all participants.

C. Participant Perceptions and Feelings towards Robot

In the questionnaire given after the last exercise session, the participants were asked to rate their perception of robot's

intelligence, and helpfulness in getting them to exercise. They were also asked to rate the level of importance they put on their participation in the exercise sessions with the robot, and their mood in general during the four exercise sessions. The rating scale was a five-point Likert scale, with the number 1 representing “Not at all”, 3 “Moderately”, and 5 “Very” (e.g., “Not at all helpful” and “Very helpful”). The question regarding user mood during the sessions contained a modified scale, where the mood options ranged from 1 “Irritated/frustrated”, to 5 “Happy/joyful”, with the medium range being 3 “Normal”.

The results of the survey questions regarding participant perception and feelings towards the robot exercise system are encouraging, as they indicate that the participants perceived the robot to have above moderate intelligence, and that the robot was more than moderately helpful in getting them to exercise. Furthermore, the participants reported placing an above moderate level of importance to the exercise sessions and normal-to-moderately pleased mood during the sessions, which is crucial for achieving long-term success in any socially assistive robot setting. A summary of the results is shown in Fig. 4 (b). The average rating was approximately 3.5 for each category; this is encouraging especially considering that the majority of participants rated each category with a score of 4 or above, and no participants rated any category with the lowest score of 1. In fact, only two participants rated the categories of helpfulness and importance below the moderate level of 3, and only one participant rated the categories of intelligence and mood below the moderate level.

Another interesting result from the study is that the participants largely (9 out of 11) favored the Workout game over the others, as shown in Fig. 4 (c). The participants (8 out of 11) also reported that they considered the robot to be more like an exercise instructor than a game conductor, companion, or none of these. Both of these results are important because they suggest that the participants have determined the robot to be an entity that they can trust and that is capable of helping them, rather than simply entertaining them.

VII. CONCLUSIONS

In this paper we have presented the design, implementation, and evaluation of a socially assistive robot that is capable of interacting with people and engaging them in a seated arm exercise scenario. A feasibility study was conducted with elderly participants; the results validate the effectiveness of the socially assistive robot in motivating users to perform simple physical exercises. Future work includes studying the role of physical embodiment by comparing the effectiveness of a physical robot to that of a computer simulation of the same robot, and also expanding the length of the study to investigate the long-term effectiveness of the socially assistive robot system.

ACKNOWLEDGMENT

We are grateful to the Southern California Presbyterian Homes staff and participating residents, with special thanks to Allen Mehta and Kerima Reed, for their help in the study preparation and execution.

REFERENCES

- [1] D. Feil-Seifer and M. J. Mataric. “Defining socially assistive robotics,” in Proc. IEEE International Conference on Rehabilitation Robotics (ICORR’05), Chicago, IL, USA, June 2005, pp. 465–468.
- [2] A. Tapus, C. Tapus, and M. J. Mataric. “The Use of Socially Assistive Robots in the Design of Intelligent Cognitive Therapies for People with Dementia”. In International Conference on Rehabilitation Robotics, Kyoto, Japan, July 2009, pp. 924-929.
- [3] J. Fasola and M. J. Mataric. “Robot Motivator: Increasing User Enjoyment and Performance on a Physical/Cognitive Task”. In Proc. IEEE International Conference on Development and Learning (ICDL), Ann Arbor, MI, USA, August 2010.
- [4] S. Dubowsky, F. Genot, S. Godding, H. Kozono, A. Skwsky, H. Yu, and L. Shen Yu. PAMM - a robotic aid to the elderly for mobility assistance and monitoring. In IEEE International Conference on Robotics and Automation, vol. 1, 570–576, Apr. 2000.
- [5] M. Montemerlo, J. Prieau, S. Thrun, and V. Varma. Experiences with a mobile robotics guide for the elderly. In Proc. of the AAAI Nat Conf on Artif Intelligence, Edmonton, Alberta, Aug 2002, pp. 587–592.
- [6] K. Wada, T. Shibata, T. Saito, and K. Tanie. Analysis of factors that bring mental effects to elderly people in robot assisted activity. In Proceedings of the International Conference on Intelligent Robots and Systems, vol. 2, 1152–1157, Oct. 2002.
- [7] C. Kidd, W. Taggart, and S. Turkle. A sociable robot to encourage social interaction among the elderly. In Int Conf on Robotics and Automation (ICRA), Orlando, USA, May 2006, pp. 3972-3976.
- [8] E.E. Baum, D. Jarjoura, A.E. Polen, D. Faur , G. Rutecki. Effectiveness of a group exercise program in a long-term care facility: a randomized pilot trial. J Am Med Dir Assoc. 2003 Mar -Apr; 4(2):74-80.
- [9] D. Dawe and R. Moore-Orr . Low-intensity, range of motion exercise: invaluable nursing care for elderly patients. Journal of Advanced Nursing, 1995. 21: 675-8.
- [10] M.D. McMurdo and L.M. Rennie. A controlled trial of exercise by residents of old people’s homes. Age and Ageing, 1993. 22:11-15.
- [11] V.S. Thomas, P.A. Hageman. Can neuromuscular strength and function in people with dementia be rehabilitated using resistance-exercise training? Results from a preliminary intervention study. J Gerontol A Biol Sci Med Sci. 2003 Aug; 58(8):746-51.
- [12] M. Zuckerman, J. Porac, D. Lathin, R. Smith, and E. L. Deci. On the importance of self-determination for intrinsically motivated behavior. Personality and Social Psychology Bulletin, 1978, 4, 443-446.
- [13] E.L. Deci, A.J. Schwartz, L. Sheinman, and R.M. Ryan. An instrument to assess adults’ orientations toward control versus autonomy with children: Reflections on intrinsic motivation and perceived competence. J of Educ Psychology, 1981, 73, 642-650.
- [14] O.C. Jenkins, C. Chu, and M. J. Mataric. “Nonlinear Spherical Shells for Approximate Principal Curves Skeletonization”. University of Southern California Center for Robotics and Embedded Systems Technical Report, CRES-04-004, 2004.
- [15] H. Fujiyoshi and A. Lipton. Real-time human motion analysis by image skeletonization. Proc. of the Workshop on Applications of Computer Vision, Princeton, NJ, USA, Oct, 1998, pp. 15-21.
- [16] S. Waldherr, S. Thrun, R. Romero, and D. Margaritis. Template-based recognition of pose and motion gestures on a mobile robot. In Proc. of the Nat Conf on Art Intel (AAAI), Madison, USA, 1998, pp. 977-982.
- [17] C. R. Wren, A. Azarbayejani, T. Darrell, and A. P. Pentland. Pfunder: Realtime tracking of the human body. IEEE Transactions on Pattern Analysis and Machine Intelligence, 19(7):780–785, 1997.
- [18] T.W. Bickmore and R.W. Picard. Establishing and maintaining long-term human-computer relationships. ACM Trans. Comput.-Hum. Interact. 12, 2 (Jun. 2005), 293-327.
- [19] NeoSpeech Text-to-Speech. 2009. <http://www.neospeech.com>