2017 International Conference on Rehabilitation Robotics (ICORR)
QEII Centre, London, UK, July 17-20, 2017.

A Learning-based Agent for Home Neurorehabilitation

Andreas Lydakis[1], Yuanliang Meng[2], Christopher Munroe[2], Yi-Ning Wu[3], Momotaz Begum[1]

Abstract— **This paper presents the iterative development of an artificially intelligent system to promote home-based neurorehabilitation. Although proper, structured practice of rehabilitation exercises at home is the key to successful recovery of motor functions, there is no home-program out there which can monitor a patient's exercise-related activities and provide corrective feedback in real time. To this end, we designed a Learning from Demonstration (LfD) based home-rehabilitation framework that combines advanced robot learning algorithms with commercially available wearable technologies. The proposed system uses exercise-related motion information and electromyography signals (EMG) of a patient to train a Markov Decision Process (MDP). The trained MDP model can enable an agent to serve as a coach for a patient. On a system level, this is the first initiative, to the best of our knowledge, to employ LfD in an health-care application to enable lay users to program an intelligent system. From a rehabilitation research perspective, this is a completely novel initiative to employ machine learning to provide interactive corrective feedback to a patient in home settings.**

I. INTRODUCTION

Home systems for neurorehabilitation are currently developed around the concept of serious gaming [14]. The key idea is to enable a patient to interact with a (virtual or augmented) game environment using motor movements that are required for rehabilitation. Custom-designed games played using Nintendo Wii [7], Sony PlayStation [10] and the Microsoft Kinect [11] are the state of the art in technology-based home rehabilitation. Standard game controllers provide limited flexibility when modifying the underlying control mechanism to suit patient-specific needs, and often exhibit poor motion tracking accuracy [16]. In recent years Microsoft Kinect has gained huge popularity in game-based motor rehabilitation research [6].

The Kinect-based systems, however, have their own limitations including the requirement of a well-designed workspace [8] as well as the occlusion problem during rehabilitation exercises [13]. In addition to these, a common pitfall of the current game-based approaches to motor rehabilitation is the inability to monitor and measure a patient's progress and adherence to the rehabilitation regimen [16]. For example, none of the existing game-based systems have the ability to detect and correct compensatory movements that a child may use while practicing at home [5]. Another limitation of these systems is their dependence on the developer. Each program generally comes with a preset range of options that can only be modified by the developers in order to incorporate any unique, patient-specific needs. This may greatly hinder the mass deployment of game-based systems at clinics and homes.

Advanced machine learning algorithms can fill these gaps by creating intelligent agents which can communicate with lay users to learn rehabilitation exercises and, later, provide interactive guidance to a patient in real-time. To this end we design an intelligent home rehabilitation (IHR) system which will bring the skills of a therapist to a patient's home through the use of interactive machine learning

Rehabilitation exercises are essentially a sequence of structured motions of different parts of the body and learning motion sequences from demonstration data is an active research domain in robotics. A large body of research in robot learning and manipulation is dedicated to designing algorithms that can extract policies (a mapping between observations and actions) from motions demonstrated by humans so that robots can mimic those human movements [2]. This approach is known as learning from demonstration (LfD). LfD is a framework for enabling an artificial agent to learn new tasks from examples provided by a teacher.

LfD [1] has gained popularity in Human-Robot Interaction research, based on its premise that, one can teach a robot new skills simply by demonstrating these skills [12] without requiring knowledge of robotics or programming. However, almost all of the contemporary LfD algorithms perform well only in scenarios where experts strictly supervise the entire operation of learning and demonstration. The presented framework aims to place LfD powered intelligent systems (e.g virtual agents, robots) in unsupervised non-laboratory settings such as in the clinic or in a patient's home.

This paper presents the development of the IHR. As the IHR is a highly patient-centered technology, the needs and expectations of the end-users (in this case, therapists and patients) are required to be reflected in the design of the algorithm and user interfaces. Accordingly, we observe an iterative process in the design of the IHR.

At system level, the paper presents a novel attempt to employ LfD in a health-care application where lay users can program and evaluate an LfD-powered intelligent system. From a rehabilitation research perspective, this is a completely novel initiative to employ machine learning to provide interactive corrective feedback to a patient in their home environment.

[1]Department of Computer Science, University of New Hampshire, Durham, NH 03824, USA `al1076, mbegum at cs.unh.edu`
[2]Department of Computer Science, University of Massachusetts Lowell, Lowell, MA 01854-2874, USA `ymeng,cmunroe at uml.edu`
[3]Department of Physical Therapy, University of Massachusetts Lowell, Lowell, MA 01854-2874, USA `YiNing_Wu at uml.edu`

II. A LfD Based Intelligent Home Rehabilitation System: The Iterative development process

The Intelligent Home Rehabilitation (IHR) system uses two commercially available technologies: two Myo armbands and a pair of augmented reality (AR) eyeglasses (Fig. 1) or another Android based device, such as a tablet. The system is expected to work in the following way:

- During a clinic visit, a therapist will demonstrates new exercises to a patient
- The wearable components of the IHR system gather and process the user's motion and electromyography (EMG) data as they perform the exercise and train a Markov Decision Process (MDP) model.
- The MDP model will drive the decision process of a 3D avatar that will appear on the display of the augmented reality eyeglasses. At home, the patient will use the same sensors and an assigned pair of AR glasses to practice the exercises. The 3D avatar will monitor the patient's practice and will provide interactive guidance when required.

To this date, the IHR system has gone through two design iterations: 1) An initial prototype was developed and evaluated by (technical) experts for technical feasibility and 2) The initial prototype was retrofitted based on the suggestions from the technical experts, and evaluated by lay users. The following sections describe these two development cycles. The research performed during the second development cycle is the contribution of this paper.

A. First Development Cycle

During the first phase we designed an IHR system which can 1) learn a motion sequence from multiple demonstrations while analyzing motion and EMG signals and execute it through a 3D avatar 2) track the activity of a person, identify when and, how they deviate from the learned exercise and provide corrective feedback. The technical feasibility of the system was evaluated through an Institutional Review Board (IRB) approved study [17]. For the sake of continuity we briefly describe the first prototype and its evaluation by technical experts.

1) System Description: While wearing an armband, which streams Inertial Measurement Unit (IMU) data and EMG signals, a user mimics a demonstrator as he performs a rehabilitation exercise. After multiple demonstrations the feature vectors corresponding to the filtered and downsampled motion and EMG data are considered as system states and are clustered using the K-means algorithm. The clustered feature vectors are also used to train a K-nearest-neighbor (KNN) classifier that classifies each incoming feature vector (during the user's practice at home) into one of the existing labels. The system also generates a score based on the accumulated reward assigned by the MDP from different state transitions.

When a user practices with the trained system at home, a 3D avatar demonstrates the entire exercise and encourages the user to perform it on their own. The avatar tracks the user's movements and, based on the training data, displays the progress the user has made toward correctly performing the exercise, the avatar, using the trained MDP model, can also provide corrective feedback by repeating the exercise from the last correct point the patient had reached. Technical details are available in [17].

2) Training and Evaluation by Technical Experts: A study was conducted to investigate the technical feasibility of the system. Three participants, all familiar with robotics but not the developed system, tried the system in a lab-based setting and provided their feedback. The scores of the different participants were analyzed along with video data to investigate if they reflected their actual performance. Participants were asked then to freely comment on how helpful they considered the system to be, how the system could be improved, and whether they would recommend it to friends or family. Some basic fundamental conclusions were derived from their answers.

- MDP model: The MDP model tracks states based on data from only the current time instant and, therefore, fails to deal with exercises that have repetitive sequences such as drawing an '8' shape. Inclusion of state history is required to enable the model to identify repetitive actions.
- 3D avatar: The mapping of motions between the human and the avatar needs to be accurate, to avoid confusion.
- Score: The evaluation of user performance needs to be more intuitive and should reflect the improvement of a user from a clinical standpoint.
- User interface: Lay users cannot deal with complex programs and prefer simple and intuitive interfaces.
- Utility: If technical issues can be resolved, the idea of helping a patient follow a rehabilitation regimen through an intelligent agent has a huge potential to be embraced.

B. Second Development Cycle

Based on these results, we retrofitted the design of the IHR system. The main changes are summarized in this section.

- Tracking of a user's current exercise state now considers a history of previous states. This modification enables the IHR to provide corrective prompts in a more accurate way for repetitive exercises. In the redesigned system, states are still assumed to be fully observable but at a given time, similar to a partially observable Markov decision process (POMDP), we evaluate the most likely system states based on a N-gram based analysis of previous states. Unlike POMDPs, however, we do not maintain a belief over all system states.
- The system now uses two armbands: one on the forearm and the other on the upper arm. This improves the mapping between arm motion and the 3D avatar. The system still runs in real time.
- A score function that better reflects the user's performance has been designed.
- Two interfaces, one for the demonstrator/therapist and the other for the patient, have been designed. Speech based interaction was added to accommodate users with limited upper limb mobility.

- Finally, the redesigned prototype has been tested with lay users: a physical therapist (a co-author of this paper) as a demonstrator and three clinical science students as patients. Each patient used the system at home for three days without any supervision or assistance from the system developers.

In the following sections we present the current IHR prototype.

1) System Description: The system was implemented in the Robot Operating System (ROS). Each Myo armband (Fig. 1) is equipped with a 9 degree-of-freedom inertial measurement unit (IMU) and 8-channel EMG sensors. The 3D avatar that acts as a virtual therapist can be projected either on the display of a tablet or an R-7 AR eyeglasses (Fig. 1). The function of the system is separated in two stages: *Training* and *Coaching*. A graphical outline can be seen in figure 1.

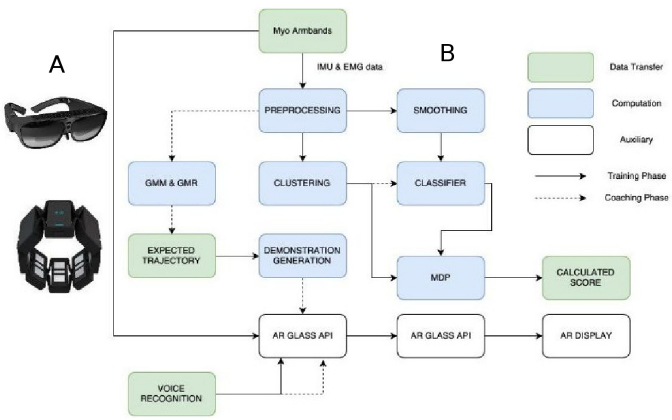

Fig. 1. A: The wearable components of the system, the R7 AR glasses (top) and the Myo armband (bottom). B: Outline of the IHR framework

Training: During this phase the therapist guides a patient to perform the exercises while tagging the critical points of the exercise using a voice recognition interface. Simple speech-based interaction enables the therapist to teach the IHR system different exercises that are tailored to each patient's need without having to undergo a specialized training process to use the system.

Data Pre-processing: The IMU and EMG data are filtered and normalized. In the case of multiple demonstrations of the same exercise, the signals are aligned using dynamic time wrapping [4]. A 5-variate Gaussian Mixture Model (GMM) is generated to model the trajectories [3]. Using Gaussian mixture regression a continuous trajectory is displayed by the 3D avatar. We use the motion and EMG data to represent state components. The state vector is defined as :

$$X_{n\times 36} = [\tilde{\alpha} EMG_{n\times 16} IMU_{n\times 20}] \tag{1}$$

Where:

- n is the number of data points
- $\tilde{\alpha}$ is a weight factor for the EMG signals (default: 1)

These vectors are used as a component to create "critical points" which will be described in the following section.

Critical Points: Critical points (C.Ps) are defined as the points significant for the exercise, for example, the endpoints for the different parts of the exercise. During a demonstration the therapist uses pre-designated voice commands to signify when a C.P. is reached. After the exercise demonstration has been completed, K-means clustering is used to identify clusters of C.Ps that represent the states. If the tagging process has produced N tags, we iterated from $0.5N - 1.5N$ (in the previous version the total time T of the exercise was used instead of N) and use the Silhouette score [15] to determine the best number of clusters, and train a KNN classifier to assign labels to incoming data vectors. These are also the points the system uses to learn a task and evaluate the performance of the patient.

Given a series of C.Ps $\{O_1, ..., O_i\}$, C.P. O_{i+1} can be determined by an n-gram model generated by the training data. N-grams are usually applied in Natural Language Processing where each of the N components could correspond to thousands of words. In such a case even a small N-gram ($\approx$ 7) would be sparse enough to significantly slow down computation. In this case however, because we usually have a small number of classes of C.Ps ($\lesssim 20$) we can allow states being represented as a tuples of six C.Ps, where the last item represents the last critical point reached.

In the case where there are not six available C.Ps, either because the exercise is too short, or because we are representing the starting portion of the exercise, we use a "wild-card" entry as the first item, that means that any entry in that position will be counted, treating the 6-gram as a 5-gram. Special "start" and "end" entries are also used. In summary, the training phase consists of clustering the data, constructing the classifier and building the n-gram from the clustered labels.

Exercise Model Creation: Using the n-gram model we would use critical points $\{O_1, ..., O_i\}$ to derive point O_{i+1}. However the patient often fails to recreate the exercise faithfully, which means that the C.Ps produced do not fit into any type available from the training data. In such a case we could use point O_{i-5} as a wild card and use the last 5 points to make the prediction. This would be equivalent to summing over all values of O_{i-5} as:

$$\begin{aligned} \tilde{O}_{i+1} &= argmax_{O_{i+1}} P(O_{i+1}|O_{i-4}O_{i-3}...O_i) \\ &= argmax_{O_{i+1}} \sum_{O_{i-5}} P(O_{i-5}) P(O_{i+1}|(O_{i-5}O_{i-4}...O_i) \end{aligned} \tag{2}$$

However, in practice it proves more efficient to count all the n-grams derived from the training data instead of computing the marginal distribution from 7-grams. This can be done using equation 3. Algorithm 1 describes the prediction process for the next C.P.

$$\begin{aligned} &\tilde{P}(O_{i+1}|(O_{i-m}O_{i-m+1}...O_i) \\ &= \frac{count(O_{i-m}O_{i-m+1}...O_iO_{i+1})}{count(O_{i-m}O_{i-m+1}...O_i)} \end{aligned} \tag{3}$$

Using six C.P.s to represent states our system can reach the same C.P. multiple times. In reality n-grams are discretized, compact representations of trajectories, with each C.P. representing a sample of the trajectory. In the current implemen-

Algorithm 1 Predict Next Critical Point

```
procedure NEXTPOINTPREDICT(M, S)
    M : the set of critical points created during training
    S = [O_{i-5}, O_{i-4}, ..., O_i]
    while S is not empty do
        if S ∉ M then                    ▷ M maps S to O_{i+1}
            S.pop()                      ▷ Pop first item
        else
            return M[S]                  ▷ O_{i+1} = M[S]
        end if
    end while
    return None
end procedure
```

tation the sample points are categorized and mapped from left to right. A more generalized representation would be:

$$\tilde{Seg}(t, t+1) = M[argmin_{S \in M.keys} Dist(seg(t-m, t), S] \tag{4}$$

Where:

- $Seg(t_1, t_2)$ is the segment of the trajectory between t_1 and t_2.
- $Dist(S_1, S_2)$ is the distance between two trajectory segments.
- $M[S]$ is the mapping of a segment S to its most likely successive segment for the time period $(t, t+1)$

When formulating the system as an MDP we use the IMU and EMG data as the MDP actions. By using training data to learn the next most likely C.P. we also learn the transitional probabilities used in the MDP. The reward function used is:

$$R_a(s, s') = P(s|s') = P(O_{i-5})P(O_{i+1}|(O_{i-5}O_{i-4}...O_i) \tag{5}$$

This reduces the model to the n-gram model based on C.Ps. The action returned by the MDP policy will be same with the most likely C.P. predicted by the n-gram model. By correctly identifying the patient's action and C.P., the avatar can give a prompt if the user that has missed the correct action by performing a demonstration of the exercise segment between the current and next point.

Having trained the MDP model we can learn its policy function $\pi(s)$ that will choose the optimal action α in each state s, maximizing the expected total reward (equation 6).

$$\sum_{t=0}^{t_{max}} \gamma^t R_{at}(s_t, s_t+1) \tag{6}$$

Using the n-gram representation the model can incorporate more complex reward functions. Since the system's current goal is to mimic the demonstrations, and provide prompts in the same manner, this reward function is sufficient and fits the nature of the tested rehabilitation tasks.

In summary the MDP components of the system are :

- S**:** states represented by a tuple of six C.Ps, produced by clustering the training data
- A**:** $X_{n \times 36} = [\tilde{\alpha} EMG_{n \times 16} IMU_{n \times 20}]$, described in 1
- $P_a(s, s')$:$P(s|s') = P(O_{i-5})P(O_{i+1}|(O_{i-5}O_{i-4}...O_i)$ as described in equation (5)
- $R_a(s, s')$: same as the transitional probability
- γ : 0.9

Coaching: During this part, the patient uses the 3D avatar to practice the exercises. They select one of the available exercises and the 3D avatar demonstrates the exercise using the model that was created in the previous step(e.g. Fig. 2 shows the 3D avatar showing a sequence of learned arm movements). It is then the patient's turn to perform the exercise. As the patient moves their arm the avatar mimics the motions. During this phase, the EMG sequences of the armbands are aligned with the ideal sequences, and the best alignment is found using Dynamic Time Wrapping (DTW). The distance between the two is used as the "EMG cost". Similarly, we calculate the trajectories of orientation coordinates from the IMU sensors and calculate the "orientation cost". DTW also gives an alignment cost for each alignment that indicates the amount of modification that was necessary in the time domain to align the signals. The evaluation rewards the user for reaching the C.P.s at the same times as shown in the demonstration. Penalties are incurred for performing the exercise too quickly or too slowly. The user's score is calculated as:

Fig. 2. The 3D Avatar demonstrating how to draw a simple sequence of movements, as an example of a rehabilitation exercise, which the user has to follow.

$$\begin{aligned} score = 100exp[-\alpha(D_{emg_1} + D_{emg_2}) \\ -\beta(D_{IMU_1} + D_{IMU_2}) - \gamma \Delta T] - 8 \times N_{prompts} \end{aligned} \tag{7}$$

Where:

- D_{IMU_x} are the euclidean distances between ideal and user trajectories for IMU signals of armband x
- D_{EMG_x} are the euclidean distances between ideal trajectories and user trajectories for EMG signals of armband x
- $\Delta T = \frac{(T - T_{demo})}{T_{demo}}$ is the relative difference in exercise execution duration
- $N_{prompts}$ is the number of prompts provided in a trial

The values used for α, β, γ were $0.08, 0.25, 0.05$ respectively. These parameters encourage the user to complete the exercise as fast as possible, but note that if we use the absolute difference to calculate ΔT we encourage the user to complete the exercise in the same amount of time as the demonstration.

2) User Interfaces: Two user interfaces have been designed based on a consultation with a rehabilitation scientist. The first is the training interface which is used by a therapist at a clinic to train the IHR system and a patient to perform some rehabilitation exercises. The second is the patient interface which is used by the patient at home. The patient can interact with the 3D avatar to receive corrective prompts and assistance only through the patient interface. Both interfaces offer the provision of natural voice based communication.

Training Interface: The training interface (Fig. 3) is used by the therapist to create the models for up to three different exercises. The therapist begins by launching the Myo armbands ("Launch Myo" button). The accumulation of drift is compensated by the frequent re-calibration of the Myo armbands ("Calibrate Myo" button). The "Clear" button is used to clear training data from previous sessions. "Train Classifier" initiates the training process, where the therapist assists the patient to reach the C.P.s of each exercise. The therapist tags those C.P.s by saying a designated word. The "Begin Trial" button is used to create the exercise model as discussed previously.

Patient Interface: The patient interface (Fig. 3) is used by patients during their practice sessions. The interface is designed to be controlled through voice commands but, mouse-based control is also available. The colored indications on the right side show the status of different hardware components: green corresponding to 'connected' and red corresponding to 'not connected'. When connected, the voice commands issued are displayed below the "Speech Detected" label. By using the "Calibrate Myo" button, the user can re-calibrate the armbands. "Task x" initiates the demonstration of task number 'x'. The demonstration can be stopped with the "Skip Demo" button. Users can use the "Stop Practice" button to stop practicing and display their score and how many times they has practiced the exercise. "Home/Reset" will stop the current task and reset the exercise counter. Using the voice command "Help", a demonstration of the next segment of the current exercise is initiated.

*3) **Pilot Study**:* We conducted an IRB approved pilot study to evaluate the prototype with a group of lay users.

Participants: A rehabilitation scientist acted as the therapist. Three female undergraduate students majoring in exercise physiology (aged $18 - 22$) participated in the study as 'patients'. Each patient completed the following questionnaire consisting of five questions. The responses were recorded in a 5-point Likert scale (Agree, Somewhat Agree, Neutral, Somewhat Disagree, Disagree) to measure a participant's exposure to modern technology.

1) I like to keep up with the latest technology.
2) I like the idea of using technology to reduce my dependence on other people.
3) I feel confident that I have the ability to learn to use technology.
4) Technology makes life easy and convenient.
5) I enjoy the challenge of figuring out high tech gadgets.

The responses indicated that all participants were familiar with everyday technology and overall, had a positive attitude toward technology.

Experimental Protocol: The study consisted of two phases. During the first phase, the therapist provided hand-to-hand demonstrations of three exercises to each of the three patients using the training interface (Fig. 3). One exercise was the same for all patients while each patient was assigned two other unique exercises. Each patient was tasked with practicing these three exercises at home, twice a day for 20min-30min using the 3D avatar and the patient interface. A brief training of how to use the patient interface was also provided to each of the three patients.

During the second phase, each patient used the system at home as prescribed by the therapist. Afterwards, they were asked to answer the following questions regarding their impressions of the IHR as a program to help follow a rehabilitation regimen. Responses were recorded on the same 5 point Likert scale as the initial questionnaire.

1) It was easy to set up the hardware
2) It was simple to start up the program
3) It was intuitive to follow the program
4) I was compliant with the program
5) I felt that I was receiving rehabilitation every day
6) I enjoyed this home program
7) I feel that the 3D avatar is interactive and understands what I am doing
8) The home-program helped me to perform the exercise when I was confused about what to do next
9) I would follow the directions the home program gives me on how to perform the exercise correctly

The participants were also asked to comment on the system in their own words. Other than the initial training, participants did not receive any expert supervision or technical support for the system. The study was intentionally designed in this way in order to investigate the potential of an LfD-powered system to be operated by lay users.

III. Results and Analysis

The small number of participants did not justify statistical analysis of the survey data. Instead we analyzed the participants' responses (Fig. 4) using descriptive statistics in order to understand the lessons learned from the second development cycle.

All the participants agreed that the system was easy to use, except for the hardware setup process. With respect to clinical utility and potential benefit for therapeutic applications, all participants agreed that they would follow the directions provided by the system to follow their regimen. Two of the participants thought that the IHR system did not feel like a "virtual therapist" when they practiced at home. The inability of the system to provide physical assistance contributed to this negative attitude toward the IHR.

With respect to the technical elegance of the IHR system, all participants reported that although the avatar was interactive, it did not provide enough feedback to them to proceed when they were stuck during an exercise.

With respect to perceived enjoyment, only one patient had a clearly positive attitude towards the system (one was neutral and the third had a negative opinion). All participants commented that the IHR was not motivating or engaging enough and suggested that providing more rewards/encouragement for the patient in a personalized way and creation of an interactive game environment can significantly increase the appeal of the system. The responses emphasize the fact that despite the intelligence of

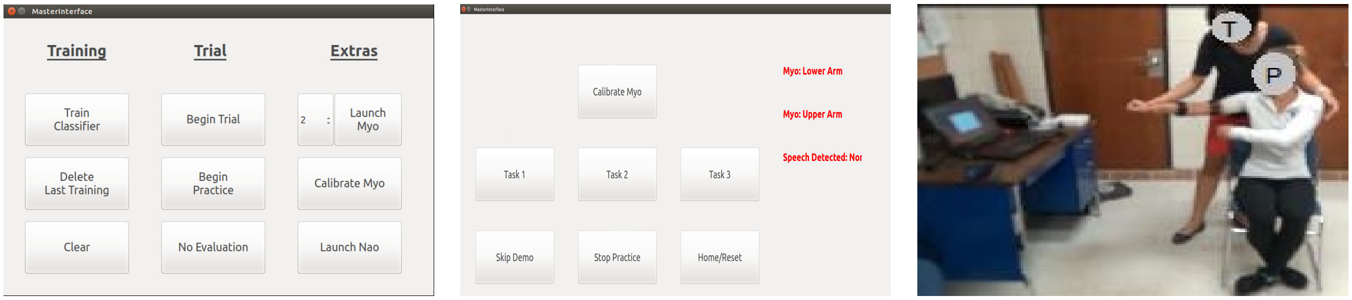

Fig. 3. Left: The trainer interface. Middle: The patient interface. Right: The therapist (labeled T) assisting one of the participant (labeled P) to perform an exercise, while using the training interface.

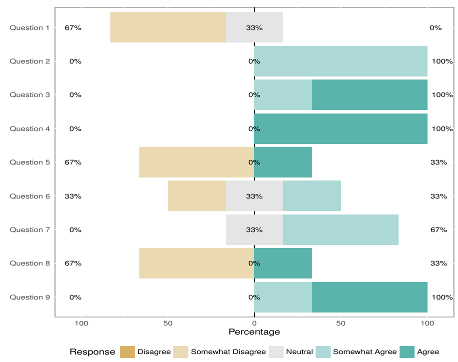

Fig. 4. Responses from the post-study questionnaire

the underlying system, faulty hardware drastically reduces the perceived overall quality of a system [9].

Finally, all participants agreed that they would recommend a similar system for home rehabilitation once all of their concerns are resolved.

IV. Conclusions and Future Work

In this paper we present the iterative development of a novel intelligent system for home-based neurorehabilitation. The presented system, inspired by the principle of LfD, enables therapists to train a system that demonstrates rehabilitation exercises based on personalized motion data and muscle activities. It also provides the patients with the ability to practice the exercises at home with expert supervision through a virtual agent.

During the second development cycle we evaluated the system with lay users. Results show that even though the intelligent system can learn arbitrary exercises from simple demonstrations and help a patient to practice at home, the user interfaces require modifications in order to be engaging and motivating.

V. Aknowledgements

This work was supported in part by the National Science Foundation (IIS 1664554). The authors would like to thank Osterhout Design Group for donating the R-7 to this project.

References

[1] C. G. Atkeson and S. Schaal. Robot learning from demonstration. In *ICML*, volume 97, pages 12–20, 1997.

[2] A. Billard and R. Siegwart. Robot learning from demonstration. *Robotics and Autonomous Systems*, 47(2):65–67, 2004.

[3] S. Calinon, F. D'halluin, E. L. Sauser, D. G. Caldwell, and A. G. Billard. Learning and reproduction of gestures by imitation. *IEEE Robotics & Automation Magazine*, 17(2):44–54, 2010.

[4] C.-Y. Chiu, S.-P. Chao, M.-Y. Wu, S.-N. Yang, and H.-C. Lin. Content-based retrieval for human motion data. *Journal of Visual Communication and Image Representation*, 15(3):446 – 466, 2004. Multimedia Database Management Systems.

[5] M. Cirstea and M. F. Levin. Compensatory strategies for reaching in stroke. *Brain*, 123(5):940–953, 2000.

[6] A. Da Gama, P. Fallavollita, V. Teichrieb, and N. Navab. Motor rehabilitation using kinect: A systematic review. *Games for health journal*, 4(2):123–135, 2015.

[7] J. E. Deutsch, M. Borbely, J. Filler, K. Huhn, and P. Guarrera-Bowlby. Use of a low-cost, commercially available gaming console (wii) for rehabilitation of an adolescent with cerebral palsy. *Physical therapy*, 88(10):1196–1207, 2008.

[8] H. Gonzalez-Jorge, B. Riveiro, E. Vazquez-Fernandez, J. Martínez-Sánchez, and P. Arias. Metrological evaluation of microsoft kinect and asus xtion sensors. *Measurement*, 46(6):1800–1806, 2013.

[9] D. Hebesberger, C. Dondrup, T. Koertner, C. Gisinger, and J. Pripfl. Lessons learned from the deployment of a long-term autonomous robot as companion in physical therapy for older adults with dementia: A mixed methods study. In *The Eleventh ACM/IEEE International Conference on Human Robot Interaction*, HRI '16, pages 27–34, Piscataway, NJ, USA, 2016. IEEE Press.

[10] M. Huber, B. Rabin, C. Docan, G. Burdea, M. E. Nwosu, M. Abdelbaky, and M. R. Golomb. Playstation 3-based tele-rehabilitation for children with hemiplegia. In *2008 Virtual Rehabilitation*, pages 105–112. IEEE, 2008.

[11] B. Lange, C.-Y. Chang, E. Suma, B. Newman, A. S. Rizzo, and M. Bolas. Development and evaluation of low cost game-based balance rehabilitation tool using the microsoft kinect sensor. In *2011 Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, pages 1831–1834. IEEE, 2011.

[12] M. N. Nicolescu and M. J. Mataric. Learning and interacting in human-robot domains. *IEEE Transactions on Systems, Man, and Cybernetics - Part A: Systems and Humans*, 31(5):419–430, Sep 2001.

[13] Š. Obdržálek, G. Kurillo, F. Ofli, R. Bajcsy, E. Seto, H. Jimison, and M. Pavel. Accuracy and robustness of kinect pose estimation in the context of coaching of elderly population. In *2012 Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, pages 1188–1193. IEEE, 2012.

[14] F. Ricciardi and L. T. D. Paolis. A comprehensive review of serious games in health professions. *Int. J. Comput. Games Technol.*, 2014:9:9–9:9, Jan. 2014.

[15] P. J. Rousseeu. Silhouettes: A graphical aid to the interpretation and validation of cluster analysis. *Journal of Computational and Applied Mathematics*, 20:53 – 65, 1987.

[16] E. Tsekleves, A. Warland, C. Kilbride, I. Paraskevopoulos, and D. Skordoulis. The use of the nintendo wii in motor rehabilitation for virtual reality interventions: a literature review. In *Virtual, Augmented Reality and Serious Games for Healthcare 1*, pages 321–344. Springer, 2014.

[17] Y. N. W. Y. Meng, C. Munroe and M. Begum. A learning from demonstration framework to promote home-based neuromotor rehabilitation. In *25th IEEE International Symposium on Robot and Human Interactive Communication*, 2016.