A Learning from Demonstration Framework to Promote Home-based Neuromotor Rehabilitation

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

Abstract— **The paper proposes a learning from demonstration (LfD) framework which will enable children with motor disabilities to perform neuromotor rehabilitation exercises at home- and community- settings. LfD, a popular robot learning paradigm, has traditionally been used to teach embodied robots different skills through demonstrations by lay users. In this paper, we propose a novel application of LfD in the health-care domain. The goal of the proposed LfD framework is to learn standard rehabilitation exercises from a therapist's demonstration during a patient's clinic visit and assist the patient to perform the exercises at home through demonstrating (using a 3D avatar) different steps of the exercise. Motion information and EMG signals of a patient are used to train a Markov Decision Process (MDP) model with different steps of the exercise from real-time demonstrations. The MDP model then tracks the progress of a patient as (s)he performs the exercise at home and provides prompts if there is any error or missed steps. The MDP model also allows quantitative evaluation of a patient's performance and improvements over time, a highly desirable property of any home-based rehabilitation system.**

I. INTRODUCTION

Learning from Demonstration (LfD) is a framework for enabling robots to learn new tasks from examples provided by a human teacher [1], [2]. The great promise of LfD is to enable lay users to teach robots new tasks simply by showing how to conduct the task and without requiring any special knowledge about robots or programming. To date LfD has been successfully used to teach robots how to assemble IKEA furniture (miniature versions) [3], how to play tic-tac-toe [4], perform pick and place tasks [5], and walking gaits [6], etc. Although the core promise of LfD is empowering lay users with the ability to control a robot, almost all of the contemporary LfD research works in scenarios where robotics experts strictly supervise the entire operation of learning and demonstration. At this stage it is unknown if robots/systems powered by LfD algorithms are actually capable of working with lay users outside of the laboratory settings.

This paper presents a novel application of LfD in the healthcare domain which will open up the possibilities of taking LfD-powered intelligent robots and systems out of the laboratory setting and enable lay users to benefit from this machine learning paradigm. The proposed application domain is neuromotor rehabilitation training for children with chronic motor disabilities. The proposed LfD framework develops a model (using Markov decision process, MDP) of a rehabilitation exercise from a therapist's real-time demonstrations during a patient's clinic visit. We use motion data (from IMUs) and EMG signals to learn the exercise sequences. Both signals (motion and EMG) are collected through a commercially available wearable sensor. The trained model then guides the patient to perform the exercise at home. Different components of the exercise are demonstrated to the patient through a 3D avatar which appears in the display screen of an augmented reality (AR) eyeglass. If the patient misses any step of the exercise, the MDP model detects that and provides necessary prompts. We have used a commercially available AR eyeglass to project the 3D avatar. Any other android devices (smart phones, tablets, etc.) can be use for this purpose. Eyeglasses, however, are hands-free and convenient to use which allows a user to perform different hand-exercises with more flexibility.

To the best of our knowledge, this is the first work that reports an application of LfD in the healthcare domain. In addition to that, no research work has been reported on the use of commercially available modern AR eyeglasses for rehabilitation services. Decentralization of therapeutic services using emerging technologies are considered as a critical need in the current health-care system [7]. The proposed LfD framework has a significant potential to promote home-based rehabilitation for individuals with motor disabilities.

II. HOME-BASED NEUROMOTOR REHABILITATION: A NEW APPLICATION DOMAIN FOR LFD

Neuromotor rehabilitation exercises refer to standard sets of muscle activities prescribed by therapists to patients with chronic motor disabilities (caused by e.g. cerebral palsy) to regain motor movements. Intensive practice of the impaired joints is thought to facilitate use-dependent neuroplasticity and promote functional recovery [8]. Unfortunately, many patients are unable to receive intensive therapy in clinical settings due to reasons including the high cost of therapy, insufficient insurance coverage, and travel inconvenience. Limited clinical visits make the expected motor improvement difficult to achieve. The use of emerging technologies - sensors, devices, robots, artificial intelligence (AI), and internet of things - for decentralization of healthcare services has become a critical need [7]. Neuromotor rehabilitation exercises are generally highly structured and repetitive in nature. Typically, it requires hand-to-hand assistance from a therapist for a patient to fully understand and master a rehabilitation exercise. For example, Fig. 1(a) shows a therapist (the third author of this paper) helping a child with cerebral palsy with a simple forearm pronation-supination

[1]Department of Computer Science, University of Massachusetts Lowell, MA 01854, USA {ymeng,cmunroe,mbegum}@cs.uml.edu
[2]Department of Physical Therapy, School of Health YiNing_Wu@uml.edu

(a) (b) (c)

Fig. 1. (a) A therapist (the third author of this paper) is helping a patient with a simple exercise (b) Myo armband (c) R-6 AR eye glasses

exercise. To achieve the desired motor recovery, it is expected that a patient will extensively practice the exercise with proper supervision/monitoring at home. This is where an intelligent system powered by LfD algorithms can contribute significantly. The LfD framework presented in this paper is expected to work in the following way:

- During a clinic visit, a therapist will demonstrate the recommended exercise to a patient.
- The patient will wear commercially available sensors that will record his/her motion data (e.g. from IMU) and EMG signals as (s)he performs the exercise with the therapist.
- All demonstration data will be batch-processed to train a MDP model with the key steps of the exercise.
- After the clinic visit, a patient will go home with a wearable sensor and a pair of AR eye glasses. At home, the MDP model, through a 3D Avatar that appears on the display screen of the AR eyeglasses, will help him/her to do the exercise and provide prompts when required.

We have developed a preliminary LfD framework using two commercially available wearable sensors: a Myo armband by Thalmic Labs Inc. (Fig. 1(b)) and an AR eyeglass, R-7, developed by Osterhout Design Group (*http://www.osterhoutgroup.com/home*) (Fig. 1(c)). The rest of the paper will discuss different components of the LfD framework.

A. System Overview

The proposed framework has two physical components: 1) a Myo armband that collects motion and EMG data and 2) a pair of AR eyeglasses to render a virtual exercise supervisor (a 3D avatar) on the display screen. Between these two, an MDP based LfD program learns tasks from demonstrations, and tracks a person's real-time progress during an exercise to provide need-based prompts.

With respect to functionality, the system has two stages: 1) *Learning*, where the system processes the demonstration data and trains a state classifier and a MDP model with the exercise sequences. An expected trajectory of the exercise is also evaluated at this stage. 2) *Coaching*, where the system performs real-time tracking of the exercise and provides prompts if any key step is missing. Figure 2 illustrates the system architecture. The system has been implemented in ROS.

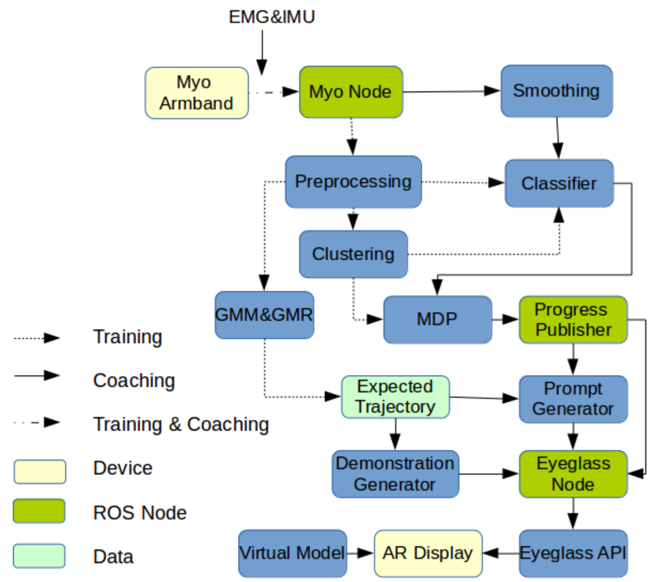

Fig. 2. An overview of the LfD framework for neuromotor rehabilitation

B. Hardware

The Myo armband (Fig. 1(b)) provides motion data from a 9 axis IMU and 8 channel EMG data to indicate muscle activation subjected to different motor activities. The R-6 (Fig. 1(c)) is a fully untethered AR eyeglass equipped with various on-board sensors (IMU, altitude sensor, light sensor, and humidity sensor). The R-6 has an auto-focus camera and a pair of stereoscopic see-through displays. Visual instructions and virtual objects can be projected in the display screen of the R-6.

C. Learning Exercises from Demonstrations

The *Learning* phase involves pre-processing and time alignment of the motion and EMG data to construct state representations, extract an expected trajectory from multiple demonstrations of the same exercise and building an MDP.

1) Data Pre-processing: An exercise is typically repeated several times (for better understanding) creating multiple data-sets for a single exercise. Increasing the number of demonstrations generally leads to a more stable system but may not be a feasible option in a real health-care setting. We have not determined the optimal number of demonstrations required to create a rich training set but found that 3 high quality demonstrations are generally sufficient to train a model. During the demonstrations of an exercise all IMU and EMG signals are recorded in a ROS bag. All data go through three pre-processing steps: down-sampling, smoothing, and normalization.

Raw signals are collected at 50 Hz. During the *Learning* phase the signals are down-sampled to 10 Hz. Rehabilitation exercises typically do not involve very fast motions and we found 10 Hz to be sufficient to train the system.

EMG data contains high-frequency noises. For example, Fig. 3(b) shows EMG signals corresponding to a simple

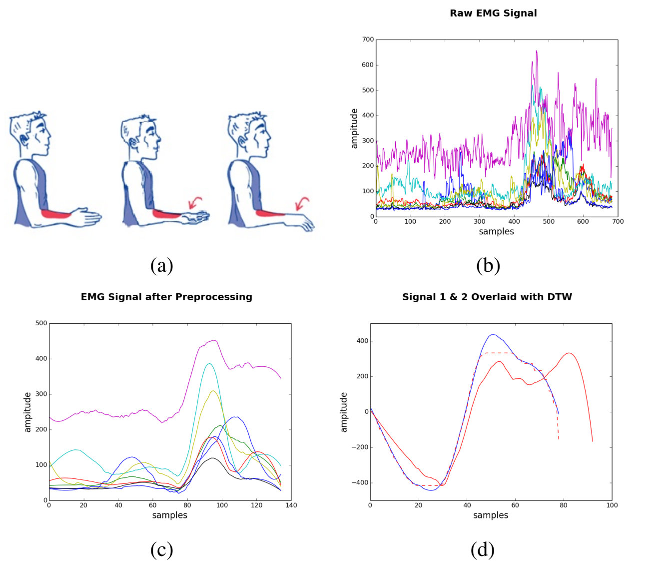

Fig. 3. Pre-processing of EMG signals corresponding to a simple forearm pronation-supination exercise shown in (a). (b) Raw EMG signals (c) Signals after down-sampling and low-pass filtering (d) DTW is applied to wrap signal 2 (red) to match with signal 1 (blue). Red dashed line shows the warped signal.

forearm pronation-supination exercise performed by an able-bodied person. A Savitzky-Golay filter [9], a computationally efficient low-pass filter, is used to smooth all signals. Fig. 3(c) shows the EMG signals after low-pass filtering. The chosen order of polynomial is 3 and the window size is 31. Note that the raw EMG signals from Myo armband return 8-bit values without having any specific unit. Some experiments show that the values are equivalent to voltages in mV amplified by a factor of 285.

For every demonstration, we collect the maximum value of the EMG signals and compute the average maximum value across trials. All values are divided by the average maximum value. Therefore, the range is mostly between $(0,1)$ while occasionally exceeding 1. Note that we obtain a single maximum value from all 8 channels, not 8 values. Accordingly, the amplitudes of some channels are always below 1. In reality, only a few EMG channels are very active with relatively high values during an exercise. The others tend to stay near zero with small changes. This is because only certain groups of muscles become active during different stages of an exercise. Similar normalization is applied to angular velocity data (from gyroscope). Angular velocities, however, are bidirectional and stay within the limit $(-1,1)$ after normalization. The absolute values of acceleration data are generally below $1g$ in typical exercises, and can also be negative. The position data from magnetometer are converted to Quaternion from Euler angles. Accordingly, Quaternion values are always within the range $(0,1)$.

2) Signal Alignment: A goal of our LfD framework is to obtain an ideal demonstration of the exercise to guide a patient. Temporal variance among signals is a critical issue when dealing with multiple demonstrations of the same exercise. Every demonstration has a different duration. If we consider a demonstration as a series of state transitions, those transitions occur at different times in different demonstrations. Therefore, all signals (EMG and motion) are aligned in the time domain using dynamic time warping (DTW) [10]. Very poorly executed demonstrations are generally discarded. Figure 3(d) shows the DTW applied to two EMG signals collected from two separate demonstrations of the same exercise shown in Fig. 3(a) (for clarity, only one dominant EMG channel is shown here). The aligned data are used in two subordinate processes: expected trajectory extraction and state representation (through clustering). The next sections will discuss the details.

3) Expected Trajectory Extraction: The proposed LfD framework learns a general representation of the task from a number of demonstrations in order to present to the user. Therefore the system must be able to provide a trajectory of coordinates as an ideal execution learned from demonstrations.

There are different ways to get the ideal trajectory from training data. We adopted the method reported in [11] where Gaussian mixture model (GMM) followed by Gaussian mixture regression (GMR) are applied on the processed demonstration set to extract a smooth trajectory to be executed by a robot.

After we align data (vectors of normalized IMU and EMG signals) from different trials, the quaternion (orientation coordinates) data are used to fit a GMM. Every quaternion coordinate is corresponding to a variable in the Gaussian model. We also add time as another variable making it a 5-variate Gaussian. Here we see the necessity of aligning signals: after adding time as a component in Gaussian, the covariance between time and other components matters. Without proper alignment, the relation between time and other components can be arbitrary.

4) State Representation: We consider motion information (angular velocity, acceleration and position) and corresponding muscle status as a "state" of the system. Trajectory coordinate variables such as joint angles have been commonly used to represent states of robot learners [11], [3]. For a rehabilitation exercise, however, it is important to include muscle activations as system state as they determine whether the exercise has been correctly executed.

As mentioned earlier, we transfers Euler angles from magnetometer to quaternion coordinates (4 dimensional vectors). As a result, IMU signals become 10 dimensional vectors. There are 8 EMG signals corresponding to 8 locations around the arm. The motion data and EMG signals are combined together to form an 18-dimensional vector to represent the "state".

$$\mathbf{X_{n\times 18}} = [\alpha \mathbf{EMG_{n\times 8}}\ \mathbf{IMU_{n\times 10}}] \tag{1}$$

Here n is the number of data points. The factor α can be used to control the weight of EMG signals with respect to motion data in state representation. For example, tasks that are more concerned about the muscle activities than the exact replication of motion can use higher values for α. We have used a default value, $\alpha = 1$.

LfD requires online learning from a limited number of demonstrations. For now the states are task-specific and need to be identified immediately after demonstrations. We have used unsupervised clustering to identify states.

K-means algorithm is used to cluster the signal vectors $\mathbf{X}$. Although we do not use time to represent state, we add time as a component to generate clusters of states here. Some tasks involve overlapping positions of arms. For example, for the task of drawing the "8" symbol, the arm will reach the midpoint twice. Although the signals, especially EMG and angular velocities, can still be significantly different in the two cases, the clustering algorithms tend to assign them the same label. This is mostly because they are more similar to each other than to other states. This is particularly true if the motion is very slow and/or has pauses which results in low amplitudes for angular velocities and EMG signals. Including time as a component of the feature vector helps to distinguish such cases during clustering. If the time stamp of data point $\mathbf{X_i}$ is denoted as t_i, the time feature can be represented as $x_t = c \times t_i$, where c is a scaling factor. In practice, we set c such that the time feature is exactly equal to the number of seconds associated with the signals. A large c value will evenly align the clusters with time intervals.

K-means algorithm requires specifying the number of states *apriori*. We iterate from 0.5T$\sim$1.5T to obtain an optimal number of states, where T is the time in second required to finish the task. In general, if arm motions are very slow, a lower number of states is needed. Then we rely on the Silhouette score [12] to pick the best clustering results. Silhouette score evaluates consistency within clusters of data, and the higher the score is, the better the clusters are formed.

After the clusters (states) S are built, we use that result to train a K-nearest neighbors (KNN) classifier to classify any new signal. The time feature was omitted while training the classifier.

The K-means and KNN algorithms are implemented with *Scikit-learn* package in Python. The GMM and GMR algorithms are adapted from the *Matlab* code provided in [11].

5) MDP Model: After obtaining states and actions from training data, we build an MDP model [13]. The goal of training a MDP model is to learn a policy function $\pi(s)$ which chooses the optimal action a in each state s. States s correspond to cluster (state) labels (obtained after applying K-means on the signal vectors) and $a \in A$ refers to cluster label of EMG patterns.The optimality is computed from the reward $R_a(s,s')$ which depends on the action a, the current state s, and/or the next state s'. Sometimes s and s' are the same, so a state transitions to itself (e.g. the case where a user pauses while performing the exercise). An optimal action maximizes the expected total reward as shown below.

$$\sum_{t=0}^{\infty} \gamma^t R_{a_t}(s_t, s_{t+1}) \tag{2}$$

$\gamma \in [0,1)$ is a discounting factor that discounts future rewards with respect to recent rewards. We have used $\gamma = 0.9$. In order to optimize the expected total reward, we need to know the transition probabilities between states conditioned on actions, as shown in equation 3.

$$P_a(s,s') = P(s_{t+1} = s' | s_t = s, a_t = a) \tag{3}$$

The frequency of the tuples (a,s,s') and (a,s) are counted in the demonstration set $\mathbf{X}$ to estimate the transition probabilities $P_a(s,s')$.

The reward function $R_a(s,s')$ assigns an immediate reward for each transition in the following way.

$$R_a(s,s') = \begin{cases} P_a(s,s') - 1 & s = s' \\ P_a(s,s') + r \times max_{s'}(EMG) & otherwise \end{cases} \tag{4}$$

Here, $max_{s'}(EMG)$ is the maximum amplitude of the normalized EMG signal in a state, obtained from the demonstration set, and r is a non-negative parameter. Such a reward function will encourage an agent to enter states that have high EMG amplitudes. States with high EMG amplitudes are generally crucial for an exercise. Sometimes, it is fine or even better to skip some uncritical states if the patient can reach the critical states from other legitimate states. The reward is set to be negative when a state transitions to itself as we want to discourage pauses or very slow motions. At any state, the probability of self-transition is actually quite high, usually much higher than the probability of transitioning to any other state.

We use the value iteration method to solve the policy. Once the policy function $a = \pi(s_t)$ is known, it can be used along with the transition probability $P(s_{t+1} = s' | s_t = s, a_t = a)$ to compute the best next states that will lead to the completion of the entire exercise. We use this information to make predictions and provide prompts to a user when (s)he fails to perform the exercise correctly. Note that the current implementation considers EMG signals simply as a state variable and does not interpreted them to meaningful real-world actions.

D. Coaching an Exercise Learned from Demonstrations

The trained MDP model is used to assist a user to perform his/her rehab exercise. The assistance is delivered through a 3D avatar that appears on the display screen of the AR eyeglasses. A user is required to wear both wearables (Myo armband and AR eyeglasses) to obtain assistance from the MDP model.

1) 3D Avatar: The 3D avatar is created with Blender (*https://www.blender.org/*) and the major functioning part is the right arm. The model is powered with jPCT (*http://www.jpct.net/*) to display motions. There are two main components: the mesh and the skeleton. The mesh is a collection of vertices which make up the shape of the model. The skeleton can be used to manipulate the mesh. Each vertex in the mesh is related to at least one bone by a scalar quantity, weight. The attribute describes how much that vertex will follow that particular bone.

The 3D model API requires rotational matrices as inputs and we use the quaternion coordinates calculated from the

(a) (b) (c) (d) (e)

Fig. 4. Demonstration of an exercise by the avatar learned from several demonstrations of a teacher. (a) A teacher demonstrates the task to a user where the user essentially follows the motion of the teacher. The user wears only the myo armband during demonstrations (b)∼(e) the avatar shows the task learned from multiple demonstrations

Euler angles of the IMU data for that purpose. Due to the limitations of the skeleton model, we employed some techniques on an ad-hoc basis to render a desirable visual effect.

2) State Classification: The trained KNN classifier classifies input signals (EMG and motion data) as a user performs the exercise. Low-pass filtering is applied on the input signals to eliminate noise. However, unlike in training data which was batch-processed, in this case we evaluate the moving average of the current and 9 previous data points. With a data frequency of 50 Hz, this results in averaging the signals over a duration of 0.2 sec. This is an acceptable time window for rehabilitation exercises which are typically performed in a slow manner.

3) Demonstration and Prompts for the User: When a user wants to do an exercise, the 3D avatar delivers a full demonstration of how to perform the exercise. For example, Fig. 4(b∼e) shows different stages of an exercise where a patient would draw the symbol 'O' in the air. Based on the GMM-GMR based model for trajectory extraction, a node on ROS publishes all expected quaternion coordinates and the Android program on the eyeglass subscribes to the message and uses it to drive the virtual arm. After demonstrating

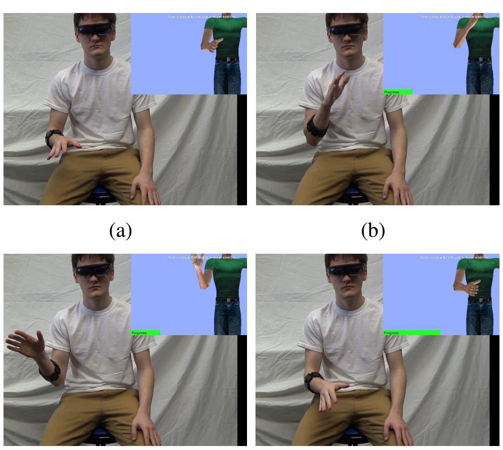

(a) (b)

(c) (d)

Fig. 5. Performing the exercise with the help of the 3D avatoar. (a)∼(e) The user (the second author of this paper) wears the Myo armband and the eyeglass to perform the task, and the AR glass display shows the progress

the task, a text is displayed in the screen asking the user to begin his/her performance. The Myo armband sends EMG and motion signals via ROS nodes at 50Hz and the proposed framework processes them in real-time to drive the virtual arm to mimic the user's arm movements. Correct state transition (within the MDP model) is considered as a progress in the path of completing the exercise. The progress is displayed as a moving bar (Figure 5). If, at some point while performing the exercise, the user forgets the correct sequence, the avatar prompts the user and shows the correct sequences to follow to complete the entire exercise. A video submitted with the paper (*ROMAN_video.mp4*) demonstrates the complete process of *Learning* and *Coaching*.

4) Performance Evaluation: A score is calculated using the reward function of the MDP model as in (5) as a measure of a user's performance. As shown in equation 4, the MDP model assigns a reward (which can be negative) for each state transition. After the task is completed, a total reward R is computed by summing up the reward correspnding to each transition. A high total reward describes a good performance. This total reward is then compared with R_0, obtained the same way from the demonstrations (training data) to compute the score.

$$score = 100 \times e^{\frac{(R-R_0)}{200}} \tag{5}$$

The total time to complete the exercise and the total number of prompts required are also reported as a measure of performance. In order to make a fair comparison, we down-sample and smooth the data collected to evaluate a patient's performance exactly the same way we do with the demonstration data. Theoretically, a performance score could be over 100, which indicates a performance is better than the demonstrations used to train the model.

III. Experiments and Results

We conducted a user study to do a preliminary evaluation of our proposed LfD framework with able-bodied participants. The study was approved by the Institutional Review Board (IRB) of the University of Massachusetts Lowell. We recruited 3 participants all of whom have some knowledge about robotics and intelligent systems. Participants, however, were not familiar with the proposed framework. Each study lasted for approximately 30 minutes and the participants were given an Amazon gift card (worth USD 5.00) as a compensation for their time. During the study each participant performed a simple exercise of drawing the symbol 'O' in the air followed by forearm pronation-supination with the help of

TABLE I

TEST DATA FROM THREE USERS

User # 1				User # 2				User # 3			
Trial No.	Time (sec)	Prompts	Score	Trial No.	Time (sec)	Prompts	Score	Trial No.	Time (sec)	Prompts	Score
1	34	0	85	1	15	0	69	1	13	0	79
2	42	2	27	2	12	0	78	2	33	2	73
3	21	1	82	3	21	1	61	3	37	1	59

the proposed LfD framework. At the beginning of the study, each participant performed the exercise while mimicking the arm motion of a researcher (the first author of this paper) for three times. The participants were asked to wear the Myo armband during this time. The demonstration set created in this way was used to train the MDP model for each specific participant. After that each participant was asked to perform the exercise without the help of the human instructor. During this time the participants were wearing the Myo armband and the AR eyeglass. Each participant performed the task three times with the help of the 3D avatar. We intentionally asked each participant to pause or move incorrectly in order to obtain their opinion about the *Coaching* behavior of the LfD framework. The study was video recorded to analyze the performance score with respect to the ground truth. Participants were asked to fill out a questionnaire consisting of three simple questions:

1. Was the 3D avatar helpful for completing the exercise?

2. What are the ways the 3D avatar could be made more helpful for you?

2. Will you recommend such a system to a friend or a family member who needs frequent assistance while performing a new task?

Table I summarizes different performance metrics. Although the task was the same, different participants generated different training data depending on their physical characteristics. In general, scores were high when a participant required less time to perform the exercise and required a fewer number of prompts. Analysis of the video shows that fewer prompts and lesser time were representative of the cases where a participant remembered the exercise and executed different steps correctly. Each incorrect step by the participant led to at least one prompt and incurred more time to complete the exercise. Thus, the reward function based score of equation (5) can be used as a reasonable representation of a participant's performance. However, much improvement of the score function is required to precisely reflect a user's performance. In the questionnaire, all participants indicated that the 3D avatar was helpful but suggested that the mapping of motion between the human and the avatar could be further improved. All participants also indicated that they would recommend such a system to those who need frequent assistance.

IV. CONCLUSION AND FUTURE WORK

We have designed a framework to apply LfD in the healthcare domain. The purpose of the proposed LfD framework is to promote home-based rehabilitation using commercially available wearable devices. Under the proposed framework an MDP model is trained with multiple demonstrations of a rehabilitation exercise by a human expert. The learned model is then used to drive a 3D avatar which can guide a user to perform the exercise in home settings. At the system level our next goal is to design a more sophisticated model (e.g. using POMDP or MOMDP) to learn complex exercises consisting of multiple steps where system states are not fully observable. With respect to utility, our next goal is to invite individuals with motor disabilities to test the proposed system, obtain their feedback and retrofit the overall design of the system.

ACKNOWLEDGMENTS

This work was supported in part by the National Science Foundation (IIS 1464226 and IIS 1552228). The authors would like to thank ODG for donating the R-6 to this project.

REFERENCES

[1] S. Chernova and A. Thomaz, *Robot Learning from Human Teachers*. Morgan & Claypool, 2014.

[2] B. Argall, S. Chernova, M. Veloso, and B. Browning, "A survey of robot learning from demonstration," *Robotics and Autonomous Systems*, vol. 57, no. 5, p. 469483, 2009.

[3] S. Niekum, S. Osentoski, G. Konidaris, S. Chitta, B. Marthi, and A. Barto, "Learning grounded finite-state representations from unstructured demonstrations," *International Journal of Robotics Research*, vol. 34, no. 2, pp. 131–157, 2015.

[4] S. Niekum, S. Osentoski, G. Konidaris, and A. G. Barto, "Learning and generalization of complex tasks from unstructured demonstrations," in *Intelligent Robots and Systems (IROS), 2012 IEEE/RSJ International Conference on*, 2012, pp. 5239–5246.

[5] M. Mhlig, M. Gienger, and J. Steil, "Interactive imitation learning of object movement skills," *Autonomous Robots*, vol. 32, no. 2, pp. 97–114, 2012.

[6] J. Nakanishi, J. Morimoto, G. Endo, G. Cheng, S. Schaal, and M. Kawato, "Learning from demonstration and adaptation of biped locomotion," *Robotics and Autonomous Systems*, vol. 47, pp. 79–91, 2004.

[7] "Report to the president: Realizing the full potential of health information technology to improve healthcare for americans: The path forward," in *http://www.whitehouse.gov/ostp/pcast*, 2010.

[8] V. W. Mark, E. Taub, and D. M. Morris, "Neuroplasticity and constraint-induced movement therapy," *Eura Medicophys*, vol. 42, no. 3, pp. 269–284, 2006.

[9] A. Savitzky and M. Golay, "Smoothing and differentiation of data by simplified least squares procedures," *Analytical Chemistry*, vol. 36, pp. 1627–1639, 1964.

[10] C.-Y. Chiu, S.-P. Chao, M.-Y. Wu, S.-N. Yang, and H.-C. Lin, "Content-based retrieval for human motion data," *Visual Communication and Image Representation*, vol. 15, pp. 446–466, 2004.

[11] S. Calinon, *Robot Programming by Demonstration: A Probabilistic Approach*. EPFL/CRC Press, 2009, ePFL Press ISBN 978-2-940222-31-5, CRC Press ISBN 978-1-4398-0867-2.

[12] P. J. Rousseeuw, "Silhouettes: a graphical aid to the interpretation and validation of cluster analysis," *Computational and Applied Mathematics*, vol. 20, pp. 53–65, 1987.

[13] M. L. Puterman, *Markov Decision Processes: Discrete Stochastic Dynamic Pro- gramming*. John Wiley & Sons, Inc., 1994.