Adaptive Therapy Strategies: Efficacy and Learning Framework

Hee-Tae Jung*, Richard G. Freedman*, Takeshi Takahashi*, Jay Ming Wong*
Shlomo Zilberstein*, Roderic A. Grupen*, and Yu-kyong Choe†
*College of Information and Computer Sciences, University of Massachusetts Amherst, Amherst, Massachusetts 01003
Email: hjung@cs.umass.edu
†Department of Communication Disorders, University of Massachusetts Amherst, Amherst, Massachusetts 01003

***Abstract*—This paper considers a data-driven framework to model target selection strategies using runtime kinematic parameters of individual patients. These models can be used to select new exercise targets that conform with the decision criteria of the therapist. We present the results from a single-subject case study with a manually written target selection function. Motivated by promising results, we propose a framework to learning customized/adaptive therapy models for individual patients. Through the data collected from a normally functioning adult, we demonstrate that it is feasible to model varying strategies from the demonstration of target selection.**

I. Introduction

In the past several decades, research interests in the use of robots in rehabilitation has increased significantly [1]. A common approach is to construct special-purpose robots and devise assistive motor controllers. They provide assistance for low functioning patients to complete otherwise unattainable movements while exercise targets are displayed on computer screens. These robots, however, rely on external software programs for generating exercise targets that could meet the needs of individual patients and adapt to the changes in the patient's motor performance.

Often exercise targets are embedded in entertaining games and motivate the patients to engage in exercise longer. In most of these approaches, patients go through different difficulty levels to essentially achieve higher performance [2]–[4]. However, many of these approaches consider generic games where exercise targets in each difficulty level are pre-programmed independent of the specific symptoms and needs of individual patients. The patient's performance is often assessed over an entire game rather than individual exercise targets. Hence, these games may neither optimally meet the needs of each individual patient nor readily adapt to the changing performance of the patients.

Addressing this, some investigate enabling therapists to have more control over customizing/adapting exercise targets. In one case, given different sets of targets, therapists can specify the thresholds and transition logic to switch among different difficulty levels [5]. However, the therapists are still forced to work with pre-determined sets of exercise targets. In another case, therapists can customize exercise targets directly, but this approach requires frequent intervention of therapists to achieve adaptation because the targets are encoded in the Cartesian space [6].

Consequently, it is necessary to validate the efficacy of the target selection approach and to investigate a means to model the therapist's intention in the selection of exercise targets through a more autonomous way of increasing the control and reducing the burden of the therapists. Section II reports the results from a single-subject case study and demonstrates that a patient can benefit even from a manually-written simple target selection function. In order to automate the process, we propose a learning framework in Section III and then introduce one way of implementing the framework in Section IV. The learned models may be used to administer therapy sessions by robots as well as computer games. Lastly, we summarize and conclude in Section V.

II. Efficacy of Adaptive Therapy Strategies

We advocate adaptive selection of individual exercise targets based on the performance of individual patients. In this section, we first validate the efficacy of adaptive therapy strategies in both standardized assessment tools, e.g. Fugl-Meyer Assessment (FMA), and in task-specific measures. A general-purpose service robot is used for the current study with residential use in mind. We adopt the inclusion criteria and therapeutic tasks from Jung et al. [6]. For the completion of the work in this paper, we reproduce them below.

A. Method

1) Participant: The post-stroke patient was recruited based on the following inclusion criteria: the participant should be at least 18 years old and had a stroke at least 6 months prior to enrollment. The assessed impairment of upper extremity motor function should score between 7 and 38 (out of 66) on the FMA. The recruited subject was a 73-year old male who experienced a stroke 10.5 years prior to enrollment and presented with hemiparesis. He scored 32 (out of 66) on the FMA at the baseline test.

2) Study Design: Fig. 1 outlines our single-subject case study design. Each condition consisted of fifteen sessions for five weeks. Sessions were performed on Mondays, Wednesdays, and Fridays unless there were scheduling conflicts. Each session consisted of three tasks which were performed for five minutes each. Between tasks, the patient took a break of approximately five minutes.

1) Task 1. The patient held two hands together and stretched his arms to reach for the robot's hand

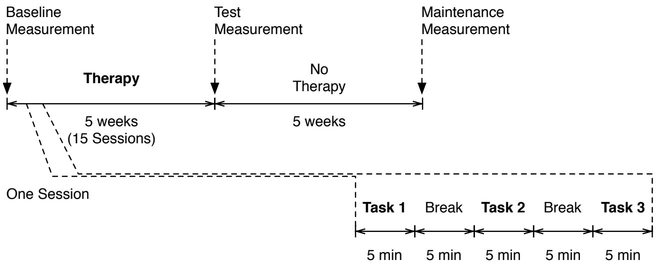

Fig. 1. The design of a single-subject case study used in this paper

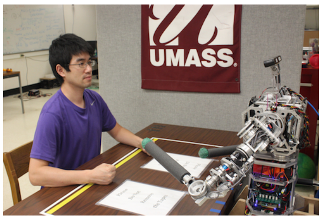

(a) Robot presenting a target

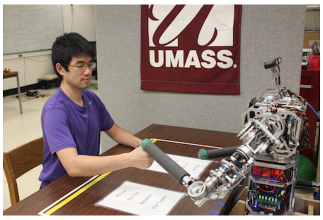

(b) Client lifting the arm

Fig. 2. A research assistant demonstrating one of the physical therapy tasks that is determined after discussion between the patient and the therapist (images reproduced from [6] with permission)

which was presented at various points within his reachable workspace. During the exercise, his intact arm assisted the impaired arm which enabled a larger range of motion.

2) Task 2. The patient lifted his impaired arm to touch the robot's hand which was presented above his hand (Fig. 2a & 2b). Since the task was challenging, we considered it successful even if he attained the presented target positions by only lifting his forearm rather than his whole arm.

3) Task 3. The patient lifted and rotated the impaired forearm to touch the robot's hand which was presented above his hand. This may appear similar to Task 2, but the recruited muscles are different.

3) Procedure: The procedure was adapted from Jung et al. [7], [8]. Before the start of the study, the patient, the therapist, and the technicians gathered along with the robot. The therapist assessed the patient's motor capability for the three tasks to determine the initial sets of target positions $(\mathbf{P}_1, \mathbf{P}_2, \mathbf{P}_3)$ as well as the arm postures that were suitable to induce desired therapeutic arm movements. The origin of the frame was placed at the patient's center of mass (COM). Since it was difficult to track the COM exactly, we approximated it by manually measuring the position of the second sacral vertebrae on the pelvis, and it was roughly controlled throughout the study by constraining the position of the chair and the sitting position of the patient.

To determine the initial set of exercise targets that is therapeutic and sufficiently challenging for the patient, the therapist controlled the positions of the robot hands to determine ideal reaching targets and asked the patient to reach and touch the robot's hand. Based on the therapist's observation and the statement of the patient, desirable targets were added to $\mathbf{P}_i$ for each task i. Each set $\mathbf{P}_i$ consisted of

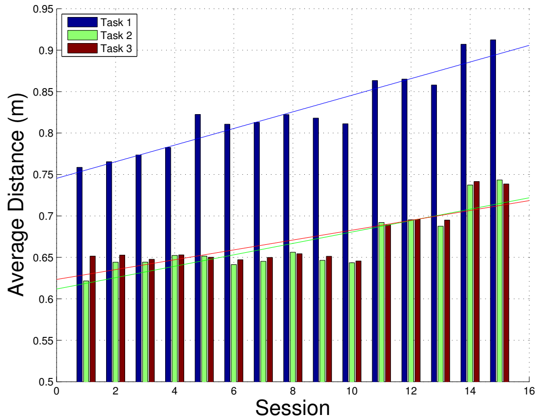

(a) The average distance of exercise targets in three tasks over fifteen sessions

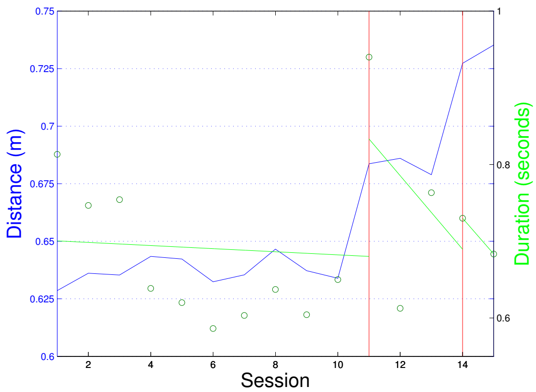

(b) The distance of one of the exercise targets (blue line, left axis) and the time taken by the patient to attain it (green circle/line, right axis) as the target position is adjusted over sessions

Fig. 3. Selected results of an exercise target over trials (best viewed in color)

N_i target positions in the Cartesian space for each task.

$$\mathbf{P}_i = \{\mathbf{p}_j | 1 \leq j \leq N_i\}$$

The therapist then instructed the technicians in the appropriate adaptation metric for each Cartesian position:

$$\mathbf{p}_j := \mathbf{p}_j + \begin{cases} \frac{\mathbf{p}_j}{|\mathbf{p}_j|}k & \text{when increasing the difficulty} \\ 0 & \text{when keeping the difficulty} \\ -\frac{\mathbf{p}_j}{2|\mathbf{p}_j|}k & \text{when reducing the difficulty} \end{cases}$$

Each target's position was moved based on the performance of the patient when reaching for it. Using the equation above, the difficulty of each target was increased when the patient successfully attained the given target position in three consecutive sessions. The difficulty was reduced when the patient failed to attain the given target position. Additionally, the patient was allowed to request to change the difficulty of each target. In this experiment, the therapist chose $k = 0.05$ (m). The robot was programmed to present the exercise targets repeatedly in sequence for the predefined duration.

4) Measures: The FMA score was used to measure the patient's overall performance change before and after the

TABLE I. THE THERAPEUTIC OUTCOMES MEASURED USING STANDARDIZED ASSESSMENT TOOLS

Tests	Baseline	Test	Maintenance
FMA	32	34	32

TABLE II. PERFORMANCE IMPROVEMENT

Task	Initial/Final Mean Distance (m)	Change (%)	t-test w/ $\alpha = 0.01$
1	0.76 / 0.91	+20	$p < 0.01$, $df = 84$
2	0.62 / 0.74	+19	$p < 0.01$, $df = 80$
3	0.65 / 0.74	+14	$p < 0.01$, $df = 90$

study. For a more detailed analysis, we employed two task-specific measures as well. First, we collected the mean distances of exercise targets in each session for each task to track the change in average distance that the patient could reach over sessions. Second, we collected the average distance of individual targets to the patient and the duration that the patient needed to reach them to track how fast the patient could attain the target over multiple sessions as its position was adjusted.

B. Results

After the completion of the five-week study, the patient scored 34 points in the FMA showing 2 points of improvement in comparison to the baseline (TABLE I). The improvement is observed in the task-specific data as well. Fig. 3a shows the overall trend in the average distance between exercise targets and the patient over sessions for all three tasks. Comparing the average distances in the first and the last sessions showed that the patient is able to reach farther, conforming to the finding in Jung et al. [8]. We ran a one-tail t-test with $\alpha = 0.01$ and observed that all the improvements are statistically significant with $p < 0.01$ for all three tasks (TABLE II). Fig. 3b shows the average distance of an individual target as well as the average time taken by the patient to achieve the target as its position is adjusted over sessions. The target position was adjusted after sessions 10 and 13, marked by the red vertical lines. Between the adjustments, a decreasing pattern in durations can be observed which conforms with the finding in Jung et al. [7].

Collectively, these results suggest that the patient can achieve significant motor performance by practicing customized and adaptive exercise targets, which suggests the importance of selecting cutomized/adaptive exercise targets. However, the model is written manually making it difficult to capture the detailed strategies of the therapist, who may consider various aspects of the patient's performance. Also, at target adaptation, each target is only projected in and out which may cover only the subspace of therapeutically meaningful space. These shortcomings are addressed below.

III. PROBLEM FORMULATION & PROPOSED FRAMEWORK

A. Problem Formulation

We assume that a therapeutic task consists of N exercise targets $\mathbf{P} = \{\mathbf{p}_1, \cdots, \mathbf{p}_N\}$ that yield observable motor performance $\mathbf{F} = \{\mathbf{f}_1, \cdots, \mathbf{f}_N\}$ of the patient where each $\mathbf{f}_n = [f_1, \cdots, f_M]^{\mathrm{T}}$ and M is the number of features describing the motor performance while the patient is attaining each target $\mathbf{p}_n$. Some examples of these tasks include:

1) a *reach-n-touch* exercise where the therapist presents a Cartesian position $\mathbf{p} = \{x, y, z\}$ and the patient reaches and touches the given target,
2) a *pick-n-place* exercise where the therapist presents an item (a salt shaker) at Cartesian position $\mathbf{p} = \{x, y, z\}$ on a table and the patient picks up the item and places it at a designated position,
3) a *ball passing* exercise where the therapist passes a ball to the patient at height and velocity $\mathbf{p} = \{h, v\}$ to perturb his balance while walking on a treadmill.

Given any of these tasks, our goal is to learn a model $\mathcal{M}$ that captures the mechanisms that are used to select the exercise targets $\mathbf{P}$ and generates new exercise targets $\mathbf{P}'$ based on the observed motor performance features $\mathbf{F}$.

B. Proposed Framework

The fundamental idea behind the proposed framework (Fig. 5) is that the therapist's decision making criteria is reflected in the patient's runtime performance $\mathbf{F}$ while the patient is attaining the given exercise targets $\mathbf{P}$.

1) Determining Exercise Targets: This first step is responsible for prescribing therapeutic tasks and selecting sufficiently challenging targets to reflect the deficits and meet the specific needs of individual patients. Naturally, it is important to have an experienced therapist assess the motor performance of the patient and discuss the potential tasks with the patient. Based on this, the therapist can prescribe a specific task and select exercise targets that are therapeutic for the individual patient, which results in $\mathbf{P}$.

2) Extracting Motor Performance Features: This step is responsible for describing the desired characteristics of each $\mathbf{p}$ by a set of features $\mathbf{f}$ that is measured while the patient attains $\mathbf{p}$. It is hypothesized that these motor performance features describe the desired characteristics of the exericse targets, e.g. the appropriate difficulty level, within the specified tasks. Typical motor performance features could be kinematic parameters, including but not limited to:

1) the time taken for the the patient to reach a target in the *reach-n-touch* exercise
2) the mean jerk of a certain joint while moving an object in the *pick-n-place* exercise
3) the peak velocity of a certain joint while receiving or throwing a ball in the *ball passing* exercise

Note that for each target that the therapist determines, there may be other exercise targets that show the same or similar motor performance. Essentially, the choice of features and the specific values will serve to define the search space from which new targets can be selected.

3) Training Models: Given the motor performance features $\mathbf{F}$ for the corresponding exercise targets $\mathbf{P}$, we need to train generative models. The choice of a model should be made to reflect the therapist's strategy and the underlying

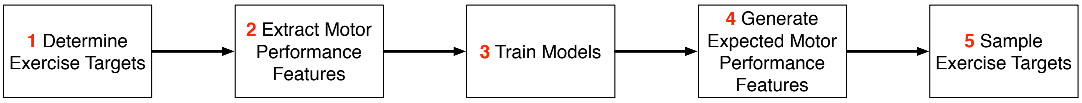

Fig. 4. The proposed framework

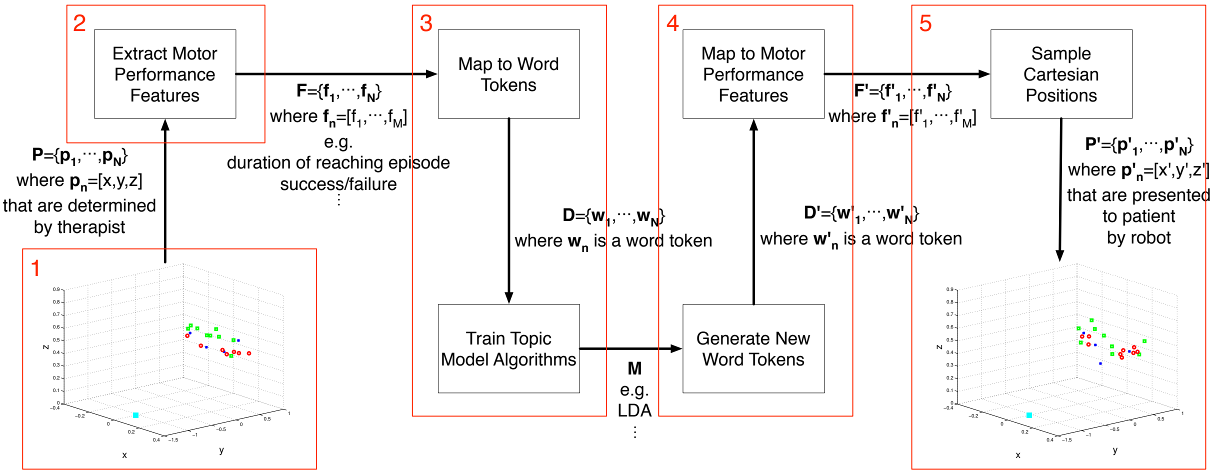

Fig. 5. The implementation of the framework in the current paper

stucture of the data. If the therapist's strategy does not change over time within each session, a simple model may suffice. However, the therapist may have a dynamic strategy that changes over time or has a particular structure. For instance, the therapist may start with easy exercise targets and select the ones that take incrementally more time until she starts regularly selecting targets of different times to take. This may be better captured by models that incorporate a temporal component.

4) Generating Motor Performance Features: We note that the models are built on the performance features $\mathbf{F}$ so that they generate new feature values $\mathbf{F}'$. Since there might be more than one target that has the generated feature values, they can be used to define the search space from which new targets can be selected based on their expected feature values, e.g. expected time to reach the target. Consequently, the following step is necessary to sample exercise targets.

5) Sampling Exercise Targets: It is most likely that the observed exercise targets in the patient's entire workspace can be sparse and we need to compute the expected motor performance at unobserved target positions before we can actually sample new targets. For example, in a *reach-n-touch* exercise, new targets should be sampled from the entire reachable workspace of the patient rather than just from a small set of reaching targets that the patient already tried. This may be achieved by applying regression techniques over the entire target space given a set of observed values. This will allow the selection of exercise targets to be from not just a set of the ones observed, but from the entire workspace.

IV. Demonstration

In this section, we apply the proposed framework to a *reach-n-touch* exercise for a normally functioning adult to demonstrate how it achieves customization/adaptation in a mock-up scenario beyond the level that a hand-built function, which was discussed in Section II, can provide.

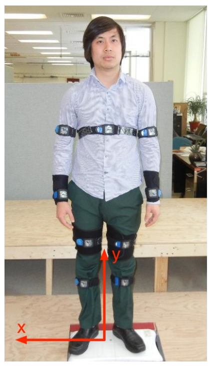

(a) On-body IMU sensors delpoyed on the subject's body

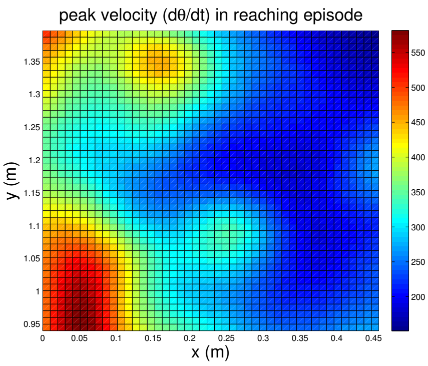

(b) Computed peak velocities during the reaching episodes; the profile plane is 30 inches away from the human subject

Fig. 6. Experiment Apparatus and a sample motor performance feature (best viewed in color)

A. Participant & Experimental Setting

Two normally functioning males participated in the data collection; one as a trainee and another as a trainer. On-body inertial measurement unit (IMU) sensors were deployed on the trainee's body to track the movement of body segments. Nine InvenSense MPU-9150 Motion Fit Wireless SDKs (0.035m$\times$0.043m$\times$0.001m) were used (InvenSense, Inc., San Jose, CA). Two sensors were placed on each limb and one was placed on the torso ($2 \times 4 + 1 = 9$). The origin of the frame was placed at the center of the trainee's feet (Fig. 6a). The trainee was instructed to hold a one-kg object in his dominant hand and keep his feet at the same positions throughout the data collection.

B. Procedure

1) Determining Exercise Targets: The trainer was provided with a graphical user interface with multiple windows, each displaying a pre-measured motor performance feature profile of the trainee on a 2-dimensional Cartesian space (Fig. 6b); the specific features provided to the user are explained below. In other words, The trainer knew what the motor performance of the trainee would look like before the selection of targets. Given this, the trainer referenced the motor performance feature of his choice and determined the exercise targets $\mathbf{P} = \{\mathbf{p}_1, \cdots, \mathbf{p}_N\}$ that would yield the performance he deemed appropriate (Fig. 7).

2) Extracting Motor Performance Features: In the literature, many investigated and used various kinematic parameters to analyze the change in the motor performance of individual patients before and after therapy given the same set of exercise targets [9]–[12]. We used these parameters to keep track of the performance of the trainee within each session and across multiple sessions.

The following parameters were computed adopting the convention in the literature. Given each $\mathbf{p} \in \mathbf{P}$, we measured a time series of joint angles $\theta_{1:T}$ for each of the nine body segments. This raw trajectory $\theta_{1:T}$ was divided into reaching, contacting, and retracting episodes. Reaching episodes were the arm movements from the resting position to the target position. Contacting episodes were the arm movements from the end of the reaching epsiode to the start of the retracting movement. Retracting episodes were the arm movements from the target position to the resting position.

Specifically, we first computed the contact point time by $t^* = \arg\max_t \theta_{1:T}$. In $[1, t^*)$, a peak velocity was computed by $v^* = \max v$ where $v = \frac{\Delta\theta}{\Delta t}$. The actual activation time t_a of the reaching movement was when the $v > 0.05v^*$ in $[1, t^*)$ for the first time. The deactivation time t_d of the reaching movement was when the $v < 0.05v^*$ in $[1, t^*)$ for the last time. $[t_a, t_d)$ defines the reaching episode. Similarly, $[t'_a, t'_d)$ was computed and represented the retracting episode. $[t_d, t'_a)$ defined the contacting episode.

For each reaching/contacting/retracting episode, we computed a peak velocity, peak acceleration, mean jerk, and movement time. Consequently, we could compute up to twelve features for each joint and a hundred and eight for all the measurable body segments $(4 * 3 * 9 = 108)$, giving us $\mathbf{F} = \{\mathbf{f}_1, \cdots, \mathbf{f}_N\}$ where $\mathbf{f}_n = \{f_1, \cdots, f_{108}\}$. For the sake of simplicity, the trainer used only a peak velocity for the right elbow joint in the reaching episode to select the targets. This work considered a family of topic models to learn and generate the underlying structure of the data. Hence, we first mapped the features to word tokens by discretizing each v^* using the step size of 8 deg/sec^2 and replaced with a corresponding symbol w_n. Given $\mathbf{F}$, we had $\mathbf{W} = \{w_1, \cdots, w_N\}$ for the training data.

3) Training Therapy Models: In this work, we explored Latent Dirichlet Allocation (LDA, [13]) and the composite model [14], which are commonly used to analyze document data[1]. LDA takes collections of sequential data, usually words, using distributions with the *bag of words* assumption which omits any sequential relationships in the given data. On the other hand, the composite model embeds LDA within a hidden Markov model (HMM) such that one state of the HMM uses LDA. This allows the model to discover and represent the structural relationships in the data.

4) Generating Motor Performance Features: After LDA and the composite model were learned, we used the learned parameters to generate new documents by sampling N words $W = \{w'_1, \cdots, w'_N\}$. Each word w'_n would correspond to some descritized peak velocity value. Since LDA treats the data as *bag of words*, the order of the words was omitted. However, the HMM states of the composite model was able to capture the sequential relationship of the data.

5) Sampling Exercise Targets: With the discretization step size used in an earlier step, we recovered the expected peak velocity for desirable exercise targets from the newly generated word token w'. Since there might exist more than a single target that could yield the same or similar value, we computed a pool of candidate targets that met $||\mu(x, y) - f'||_2 < 2\sigma(x, y)$ where μ and σ are the expected feature value and the standard deviation that we computed using Gaussian Processes with Matérn covariance function [15]. Fig. 8–10 show the subset of these pools for the five generated exercise targets in a row in the beginning and in the middle of the session.

C. Results

Here we present the results from a learned LDA model with 5 topics and a learned composite model with 5 topics and 8 states. Fig. 7–Fig. 10 show the exercise targets (white filled circles) overlapped on the peak velocity feature profile. In the training data, the trainer started with exercise targets of high peak velocities for warming-up exercise and subsequently started intensive therapy by selecting targets of low peak velocities (Fig. 7). As expected, LDA presents the combination of these targets (Fig. 8) but the composite model was able to capture the change in the trainer's strategy (Fig. 9). Fig. 10 demonstrates how the proposed framework can adapt to the change in the motor performance profile.

V. Conclusion

In this paper, we addressed the problem of adaptive therapy strategies in two parts. We first investigated the efficacy of the adaptive sampling of exercise targets and showed that it can lead to an increase in a patient's motor performance. Motivated by these results, we then proposed a data-driven framework to learn models of the therapist's target selection strategies. The realization of the proposed framework was demonstrated using LDA[2] and the composite model.

[1] Note that, however, there is no constraint in the choice of generative models and one can choose any necessary representation that corresponds to their choice of models/algorithms.

[2] More detailed results using LDA can be found in Jung et al. [16].

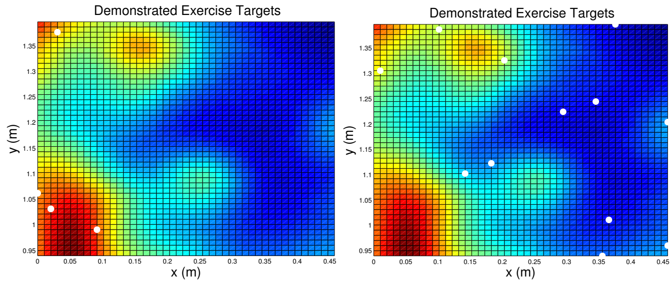

(a) (b)

Fig. 7. Demonstrated exercise targets in the beginning and the middle of the session. (best viewed in color)

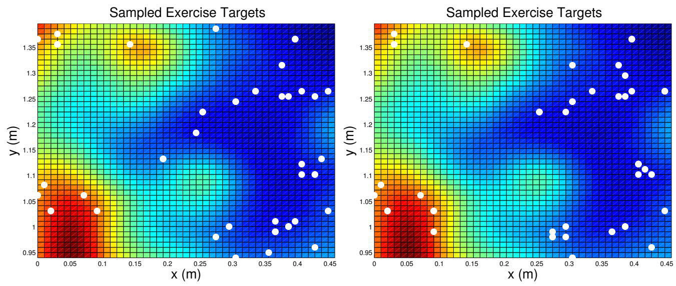

(a) (b)

Fig. 8. LDA selects exercise targets of similar motor performance features throughout an autonomous session. (best viewed in color)

ACKNOWLEDGMENT

The authors thank Jennifer Bessire and Tammie Foster for their help. This work was supported by an award from the American Heart Association (12CRP9010007).

REFERENCES

[1] G. Prange *et al.*, "Systematic review of the effect of robot-aided therapy on recovery of the hemiparetic arm after stroke," *Journal of Rehabilitation Research and Development*, vol. 43, no. 2, pp. 171–184, 2006.

[2] J. W. Burke *et al.*, "Optimising engagement for stroke rehabilitation using serious games," *Visual Computing*, vol. 25, pp. 1085–1099, 2009.

[3] M. Pirovano *et al.*, "Self-adaptive games for rehabilitation at home," in *Proceedings of the IEEE Conference on Computational Intelligence and Games*, 2012.

[4] N. Shirzad and M. V. der Loos, "Adaptation of task difficulty in rehabilitation exercises based on the user's motor performance and physiological responses," in *Proceedings of the IEEE/RAS-EMBS International Conference on Rehabilitation Robotics*, 2013.

[5] S. Gobel *et al.*, "Serious games for health – personalized exergames," in *Proceedings of the ACM International Conference on Multimedia*, 2010.

[6] H. Jung *et al.*, "Towards extended virtual presence of the therapist in stroke rehabilitation," in *Proceedings of the IEEE/RAS-EMBS International Conference on Rehabilitation Robotics*, 2013.

[7] ——, "Upper-limb exercises for stroke patients through the direct engagement of an embodied agent," in *Proceedings of the 6th ACM/IEEE International Conference on Human-Robot Interaction*, 2011.

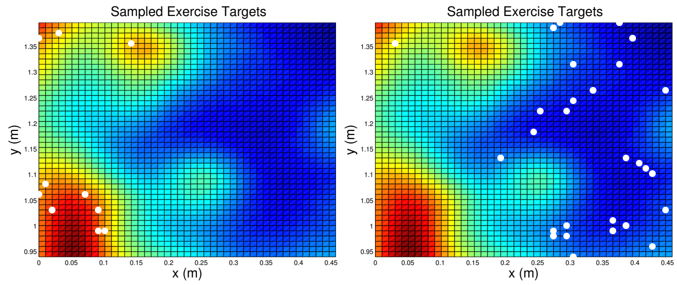

(a) (b)

Fig. 9. The composite model selects exercist targets of different features as a session progresses. (best viewed in color)

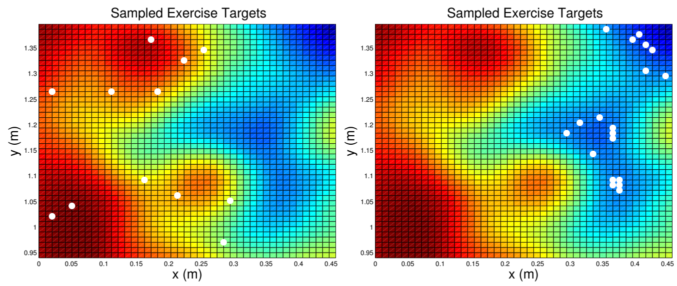

(a) (b)

Fig. 10. The targets selected by the composite model as the patient's performance changes. The proposed framework can adaptively select targets since the selection is made based on the targets' features. (best viewed in color)

[8] ——, "Upper extremity physical therapy for stroke patients using a general purpose robot," in *Proceedings of the 20th IEEE International Symposium on Robot and Human Interactive Communication*, 2011.

[9] S. E. Fasoli *et al.*, "Context and goal-directed movement: The effect of materials-based occupation," *Occupational Therapy Journal of Research: Occupation, Participation, and Health*, vol. 22, pp. 119–128, 2002.

[10] T. L. Gasser-Wieland and M. S. Rice, "Occupational embeddedness during a reaching and placing task with survivors of cerebral vascular accident," *Occupational Therapy Journal of Research: Occupation, Participation, and Health*, vol. 22, pp. 156–160, 2002.

[11] G. T. Thielman *et al.*, "Rehabilitation of reaching after stroke: Task-related training versus progressive resistive exercise," *Archives of Physical Medicine and Rehabilitation*, vol. 85, pp. 1613–1618, 2004.

[12] P. M. van Vliet and M. R. Sheridan, "Coordination between reaching and grasping in patients with hemiparesis and healthy subjects," *Archives of Physical Medicine and Rehabilitation*, vol. 88, pp. 1325–1330, 2007.

[13] D. M. Blei *et al.*, "Latent Dirichlet Allocation," *Journal of Machine Learning Research*, vol. 3, pp. 993–1022, 2003.

[14] T. L. Griffiths *et al.*, "Integrating topics and syntax." in *Proceedings of the 18th Annual Conference on Neural Information Processing Systems*, 2004.

[15] C. E. Rasmussen and C. K. I. Williams, *Gaussian Processes for Machine Learning*. Cambridge, MA: The MIT Press, 2006.

[16] H. Jung *et al.*, "Learning therapy strategies from demonstration using latent dirichlet allocation," in *Proceedings of the ACM International Conference on Intelligent User Interface*, 2015.