

Learning Therapy Strategies from Demonstration Using Latent Dirichlet Allocation

Hee-Tae Jung¹ Richard G. Freedman¹ Tammie Foster²
Yu-Kyong Choe³ Shlomo Zilberstein¹ Roderic A. Grupen¹

¹ School of Computer Science, University of Massachusetts, Amherst, MA 01003

² Center for Human Motion, Cooley Dickinson Hospital, Northampton, MA 01061

³ Department of Communication Disorders, University of Massachusetts, Amherst, MA 01003

ABSTRACT

The use of robots in stroke rehabilitation has become a popular trend in rehabilitation robotics. However, despite the acknowledged value of customized service for individual patients, research on programming adaptive therapy for individual patients has received little attention. The goal of the current study is to model teletherapy sessions in the form of a generative process for autonomous therapy that approximate the demonstrations of the therapist. The resulting autonomous programs for therapy may imitate the strategy that the therapist might have employed and reinforce therapeutic exercises between teletherapy sessions. We propose to encode the therapist's decision criteria in terms of the patient's motor performance features. Specifically, in this work, we apply Latent Dirichlet Allocation on the batch data collected during teletherapy sessions between a single stroke patient and a single therapist. Using the resulting models, the therapeutic exercise targets are generated and are verified with the same therapist who generated the data.

Author Keywords

Robot-Mediated Therapy; Customized Therapy; Learning from Demonstration; Latent Dirichlet Allocation.

ACM Classification Keywords

H.1.2. Models and Principles: User/Machine Systems; I.2.6. Artificial Intelligence: Learning; J.3 Life and Medical Sciences: Health

INTRODUCTION

Many countries are experiencing an increase in the elderly population, which may lead to spiraling healthcare costs and shortages of trained professionals to address the needs of the aging population [26]. Robots are often considered for healthcare applications including stroke rehabilitation which is the focus of the proposed work. To date, a common approach is to develop multiple special-purpose robots each of which is designed to serve a single particular rehabilitation goal. Good examples include the exoskeleton rehabilitation

robots that are built to reconstruct upper or lower limb movements [4, 11]. Due to their size, cost, and special purpose nature, these robots are most affordable when they are shared by a large number of users in hospitals or rehabilitation centers [17]. This work takes a different approach by proposing that general-purpose personal robots—those that can reside with a patient and deliver various services—can be programmed by service providers to deliver some services autonomously, in a data-driven manner that adapts to changing client needs. In both cases, it is important to encourage patients to exercise with voluntary movements and therefore necessary to provide adaptive exercise targets with appropriate difficulty levels.

Recently, researchers have started exploring ways to adapt the difficulty levels of exercise programs based on patient performance. For instance, Perry *et al.* increase the difficulty levels of exercise games when patients exceed predefined achievement thresholds [20]. Shirzad *et al.* explore various supervised learning methods to predict patients' intention to adjust the difficulty levels of games [22]. However, these approaches rely on generic games with predetermined exercise programs and difficulty levels. Also, they do not allow therapists to prescribe exercise targets or specify difficulty levels directly. Jung *et al.* allow therapists to select exercise targets to meet the specific needs of individual patients [16]. However, their approach encodes the targets in the Cartesian space and requires the frequent intervention of therapists to re-select the targets as the motor performance of patients change.

The goal of our research is to develop a new approach to create data-driven generative models that capture the decision making capacity of human therapists during therapy sessions with patients and verify that the new sessions generated by the resulting models can be vetted and endorsed by therapists. Specifically, we propose to use Latent Dirichlet Allocation (LDA) and learn the underlying features of therapist-selected targets. This may reveal the therapist's strategy as well as the motor characteristics that are specific to an individual patient. We test our approach using batch data of therapeutic exercise targets prescribed by an actual therapist for an actual stroke patient during teletherapy sessions. We contend that these technologies can be used to reinforce therapy delivered by remote service providers and may, thus, improve their productivity while it also improves client outcomes.

BACKGROUND

Robot-Mediated Therapy Using robots for stroke rehabili-

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

IUI 2015, March 29–April 1, 2015, Atlanta, GA, USA.

Copyright is held by the owner/author(s). Publication rights licensed to ACM.

ACM 978-1-4503-3306-1/15/03 ...\$15.00.

<http://dx.doi.org/10.1145/2678025.2701403>

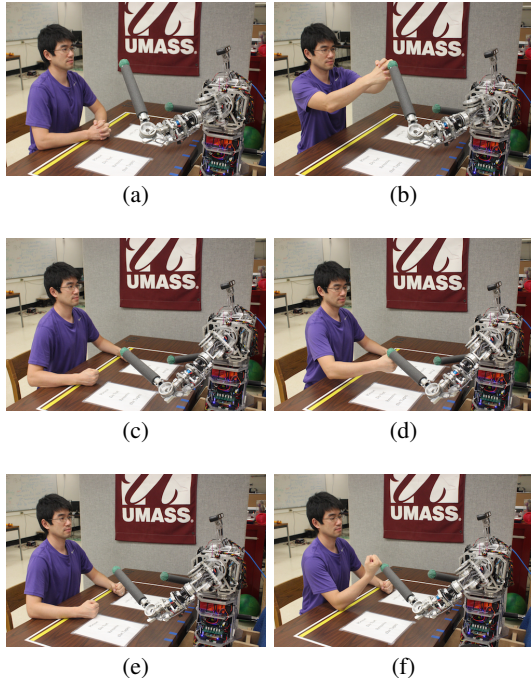


Figure 1: A research assistant demonstrating physical teletherapy sessions. The robot presents a reaching target position in 3D Cartesian space, which is determined by the therapist based on the individual client’s motor performance. In task 1, the client is instructed to hold two hands together and reach the target (a, b). In task 2, the client is instructed to lift his impaired arm to reach the target (c, d). In task 3, the client lifts and rotates his impaired forearm to reach the target (e, f).

tation has received a lot of attention. Most effort has focused on developing new robot platforms and validating their feasibility [2, 19, 21]. The rehabilitation robots realize hands-on style therapy where stroke patients either hold onto or place their impaired arm on the end-point or the exoskeletal components of these robots. During therapy sessions, these robots provide assistance for low functioning patients to complete otherwise unattainable exercise targets in pre-programmed computer games with various exercise targets [5]. Since it is important that the patient makes voluntary movements with minimum support, the patient may benefit from carefully selected exercise targets tailored to his motor impairment.

Some explore the possibility of utilizing general purpose robots. Jung et al. explore the idea by Dijkers et al. [7] and induce voluntary exercise movements from post-stroke patients in both automated and tele-operated modes [13, 16]. In this approach, the challenge level of exercise is mostly regulated by the way the exercise targets are presented while the patient attains these targets through voluntary movements. If appropriate exercise targets are provided, empirical studies report that this can lead to larger range of motion [13–15]. Hence, it is apparent that the selection of the targets need to be tailored to the impairment and meet the needs of individual patients.

Latent Dirichlet Allocation LDA [1] is an unsupervised machine learning technique commonly used in text analysis. The underlying concept behind LDA is that an observed collection of D documents contains information about T topics, distributions $\phi_{t \in \{1, \dots, T\}}$ over a list of words. However, these topics are not known *a-priori* and a single document may discuss more than one of these topics, represented by a distribution $\theta_{d \in \{1, \dots, D\}}$. The topics are concepts which may be used as keywords to summarize the collection of documents, and a bag-of-words assumption (ordering does not matter) is applied to the collection of documents to emphasize this keywords notion. Thus identifying words such as “carrot” and “orange” anywhere in a document implies that it is likely to contain a topic about food, produce, or orange edible items; the unsupervised approach requires human experts to interpret the learned topics.

LDA finds these topics by clustering words in order to maximize underlying parameters that are part of a generative process explaining how the documents were created. In summary, the θ_d and ϕ_t distributions are each drawn from respective Dirichlet distributions, and then each word w_i in document d is generated by sampling topic $z_i \sim \theta_d$ so that $w_i \sim \phi_{z_i}$. This means that the likelihood of θ , ϕ , and the latent topic assignments \mathbf{z} should be maximized for text \mathbf{w} :

$$P(\mathbf{z}, \theta, \phi | \mathbf{w}, \alpha, \beta) = \frac{P(\mathbf{w} | \mathbf{z}, \phi) P(\mathbf{z} | \theta) P(\theta | \alpha) P(\phi | \beta)}{P(\mathbf{w} | \alpha, \beta)}$$

where the denominator is a normalizing constant and α, β are hyperparameters describing the Dirichlet distributions. Once these parameters are learned using Gibbs sampling, we may use them to infer topics of newly observed documents assumed to be derived from the same generative process as well as find the likelihood that such documents were actually generated by the same model. LDA has also been widely used outside of text analysis by applying the bag-of-words assumption to other collections of objects including pixel regions [24] for semantic image analysis, streams of sensor data [8, 12, 18, 27] for activity recognition, and sequences of images [3, 25] for activity recognition and segmentation.

PROPOSED APPROACH

Learning Therapy Strategies The fundamental idea behind our approach is that the therapist’s decision making criteria is reflected in the runtime performance of the patient during the prescription of therapeutic exercises. Given therapeutic tasks, exercise targets should be prescribed in a way that challenge an individual patient to an appropriate extent. In each therapy session, let us assume that the therapist determines N exercise targets. For each described exercise target p_n , the runtime motor performance of the patient $\mathbf{x}_n = [x_1, \dots, x_M]^T$ can be measured where M is the number of features describing the motor performance while the patient is attaining the target. We hypothesize that $(\mathbf{x}_1, \dots, \mathbf{x}_N)$ can describe the dynamics of the challenges that the therapist wants to impose during the therapy sessions for the particular patient. In this work, each feature’s measurement x_m is mapped to a representative symbol and \mathbf{x}_n concatenates these M symbols to form word token w_n . Each therapy session then contains a

vector of word tokens $\mathbf{w} = [w_1, \dots, w_N]^T$ so that the session has a formalized document representation. This allows us to apply state-of-the-art topic modeling methods to capture the underlying structure of these words, i.e. the patient’s performance features. We use LDA as it is a well-studied algorithm and we are uncertain which criteria best cluster reaching targets for physical therapy sessions. We consider two performance features in this work: the success of reaching movement and the duration taken while reaching. The difficulty of each given exercise target is modeled using these features.

Generating New Exercise Targets Using a teleoperation interface, the therapist can lay out therapeutic exercises in real time which requires constant intervention. However, we can use the learned parameters to generate new exercise targets according to the criteria learned from demonstrations. This will allow the therapist to see more patients per unit of time. Because each word in a session’s document represents a single point in space, the bag-of-words assumption appears more applicable to our domain. Unlike text data in which word ordering is necessary for comprehensibility, the set of points may be reordered and yield a similar therapeutic treatment. Hence we perform the generating process from which LDA derived our training data in order to create original documents that propose a set of points for new exercise targets in future training sessions. That is, for a new session containing N' points, we sample $z'_1, \dots, z'_{N'} \sim \theta' \sim \text{Dirichlet}(\alpha)$ and then sample $w'_i \sim \phi_{z'_i}$ for each $i \in \{1, \dots, N'\}$. For each w'_i , we can approximately invert the symbolic encoding to find Cartesian position $f^{-1}(w'_i) = p'_i$. Consequently, as the patient becomes tired or improves his motor performance, the exercise target p' is adjusted to yield similar motor performance features. Due to limited space, we simply note that we use the patient’s performance database to compute f^{-1} .

EXPERIMENTS

Participants In order to validate the approach, we used a small dataset collected from five teletherapy sessions with a real patient and a real therapist. The patient was a 73-year-old male who experienced a stroke 10.5 years prior to the data collection. He scored 32 out of 66 points on the upper extremity portion of the Fugl-Meyer Assessments [6, 9]. The therapist had 22 years of experience with expertise in neurological patients including stroke patients.

Data Collection In a single day, three tasks were performed using teleoperation. The therapist remotely controlled the robot and determined exercise targets for three different tasks following the method presented in Jung et al. [16]. See Figure 1 for a detailed description of the tasks. This was repeated five times yielding fifteen sets of data (3×5).

The exercise targets and the corresponding patient’s performance were collected as batch data. For each reaching target in the Cartesian space $p \in \mathbb{R}^3$, two motor performance features were measured. The first was whether or not the patient successfully attained the target, and the second was the duration of time taken to reach the target from resting position. Note that the first feature $x_1 \in \{\text{success, failure}\}$ is discrete by nature while the second feature $x_2 \in \mathbb{R}$ is continuous and

needs to be discretized. In this work, x_2 is discretized systematically using a specified constant interval, i.e. $i = 0.3$ for task 1 and $i = 0.5$ tasks 2 and 3. The intervals were determined empirically. These feature values were mapped to the corresponding symbols and concatenated.

Therapist’s Evaluation It is important that therapists find the results useful. We investigated two criteria. First of all, the learned topics describing exercise targets should make sense, revealing some characteristics of the patient’s performance and the pattern of exercise targets. The exercise targets were shown to the therapist in the form of plots and a list of the patient’s corresponding motor performance. We asked if the therapist could identify any patterns and whether the learned topics made sense to her. Furthermore, we presented some newly generated therapy sessions and asked if she believed that they would be appropriate in terms of difficulty.

RESULTS & DISCUSSION

The model was trained using different values of T and the one providing the greatest log likelihood (LL) was chosen. For task 1, x_2 was discretized using $i = 0.3$ resulting in five intervals. $T = 2$ topics had the best LL where the first topic (red) included targets that were successfully attained in $0.6 \leq t < 1.2$ seconds by the patient. The second topic (blue) included targets successfully attained mostly within $t < 0.6$ seconds. There were a small number of targets that took $t \geq 1.2$ seconds to attain or that the patient failed to attain. When the results were presented to the therapist, she found this plausible and stated that:

He has flexor tone and it is easier [for him] to use his tone [to reach the targets]. For the external rotation, it’s much more difficult to reach. It is easier to cross the middle line and reach because he can use his torso.

For task 2, x_2 was discretized using $i = 0.5$ resulting in five intervals. The patient was instructed to make unassisted voluntary movements, taking more time in general. Based on the LL, $T = 3$ topics were learned. The first topic (red) includes targets attained in $0.5 \leq t < 1.0$ seconds, the second topic (blue) includes targets attainable in either $t < 0.5$ seconds or $1.5 \leq t < 2.5$ seconds, and the last topic (green) includes mostly targets the patient could not obtain (as well as a small number of targets attained in $1.0 \leq t < 1.5$ seconds). As seen in Figure 2b, topics were credited to the targets’ heights.

For task 3, x_2 was discretized using $i = 0.5$ resulting in four intervals. Based on the LL, $T = 2$ topics were learned. The first topic (blue) includes attainable targets in $0.5 \leq t < 1.0$ seconds while the second topic (red) includes the remaining targets. This can be interpreted as one topic including the exercise targets that the patient can attain with *some* effort. The other topic includes targets that are attainable with minimum or significant effort in addition to those not attainable at all.

The therapist was given five suggested therapy sessions per task that were generated using the trained LDA models. Between thirty and fifty targets were randomly generated per session. For task 1, she felt that four out of five closely captured the therapy she would provide and would be willing to let the robot use all five suggested sets of exercise targets. One

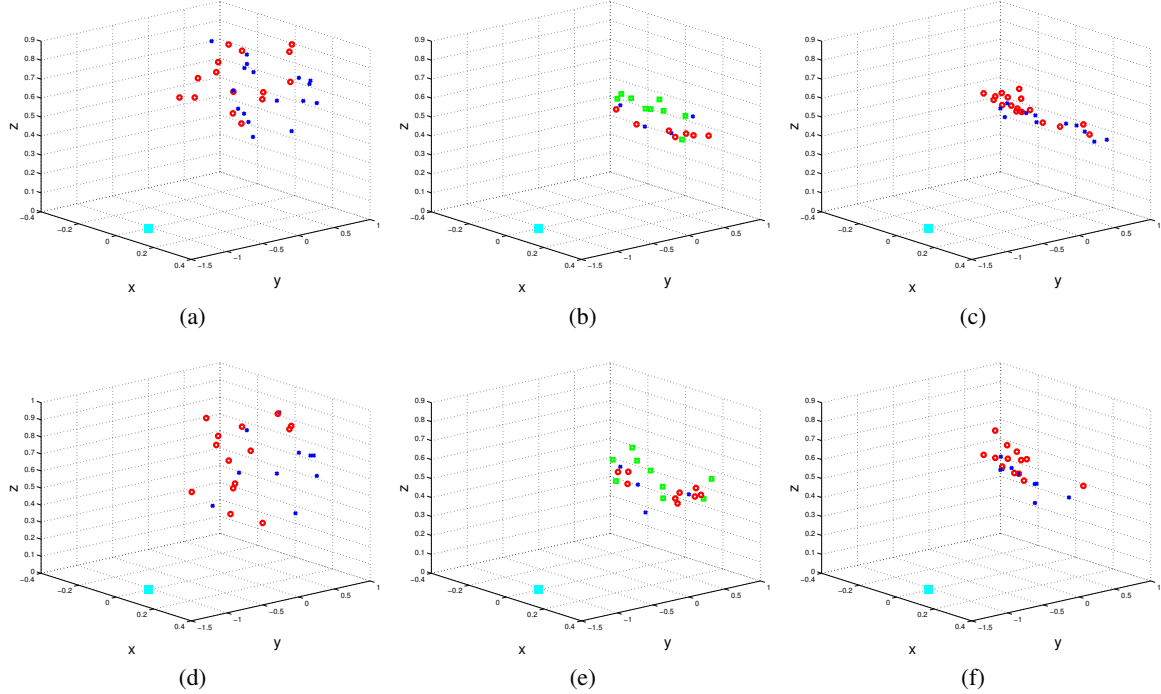


Figure 2: (a), (b), (c) are the target positions the therapist determined during her teletherapy sessions for tasks 1, 2, and 3 respectively. (d), (e), (f) are the target positions the trained model generated for autonomous therapy sessions for tasks 1, 2, and 3 respectively. A marker with cyan color at $(0, -1.05, 0)$ is the approximate sitting position of the patient. Units are in meters.

of the suggested sets of targets had ‘less motions to the left (horizontal adduction) the way that was harder for the patient to reach.’ For task 2, the therapist thought all five suggested sets captured the therapy and would let the robot use them. However, for task 3, the therapist was satisfied with only one out of the five suggested sets due to the small number of targets in the four sets. She responded she would have endorsed them if the sufficient numbers were generated.

During the interview, the therapist thought the choice of features made sense because both play an important role in how she chooses exercise targets. However, she pointed out that it would be useful to separate cognitive and motor movement delay. In this work, they were encoded as a single duration feature. Additionally, she suggested that different spatial areas may challenge the patient differently because of the flexor tone specific to the patient. Lastly, she mentioned that it might be useful if temporal patterns can be captured.

It was encouraging that she found the presented plots useful for understanding the motor performance of the patient. This implies that topic modeling may be used to both provide visual feedback and model the therapist’s strategy.

CONCLUSION

Through a single subject case study, we demonstrated that therapy strategies can be modeled in terms of the patient’s motor performance and that the models can be automatically learned using LDA. We also showed that the learned parameters can be used to generate sets of new exercise targets that

conform with the therapist’s strategies. Using the proposed approach, the robot might be able to imitate the therapist’s strategies and provide appropriately challenging targets even in his/her absence. This can increase the efficiency of the therapist’s time in practice.

The proposed approach can be improved in multiple ways. For example, it is well known that LDA needs a large amount of training data. Here, the number of features and the available values were small which might have been the reason that LDA worked using the small data set. Further investigation needs to be done to determine whether simpler models with fewer parameters, such as the Bayesian mixture model, should be used. This is especially important because the size of the data we can collect in practice is small.

As the therapist implied, it is possible for her decision making process to evolve even within the same therapy session. Since LDA treats the data as a bag-of-words, it cannot capture this temporal distinction. It might be interesting to see if this can be learned using temporal extensions of LDA which integrate hidden Markov models or n -grams to acknowledge structure in the data [10, 23]. There are many variations of LDA that can address the new demands that are likely to arise as we continue to enrich the domain.

ACKNOWLEDGMENTS

This work was supported in part by an award from the American Heart Association (12CRP9010007) and the National Science Foundation (IIS-1116917).

REFERENCES

1. Blei, D. M., Ng, A. Y., and Jordan, M. I. Latent Dirichlet Allocation. *Journal of Machine Learning Research* 3 (2003), 993–1022.
2. Brewer, B. R., McDowell, S. K., and Worthen-Chaudhari, L. C. Poststroke upper extremity rehabilitation: a review of robotic systems and clinical results. *Topics in stroke rehabilitation* 14, 6 (2007), 22–44.
3. Buchsbaum, D., Canini, K. R., and Griffiths. Segmenting and recognizing human action using low-level video features. In *Proceedings of the 33rd Annual Conference of the Cognitive Science Society* (2011).
4. Colombo, G., Wirz, M., and Dietz, V. Driven gait orthosis for improvement of locomotor training in paraplegic patients. *Spinal Cord* 39, 5 (2001), 252–255.
5. Colombo, R., Pisano, F., Mazzone, A., Delconte, C., Micera, S., Carrozza, M. C., Dario, P., and Minuco, G. Design strategies to improve patient motivation during robot-aided rehabilitation. *Journal of NeuroEngineering and Rehabilitation* 4, 3 (2007).
6. Deakin, A., Hill, H., and Pomeroy, V. M. Rough guide to the fugl-meyer assessment upper limb section. *Physiotherapy* 89, 12 (2003), 751–763.
7. Dijkers, M. P., deBear, P. C., Erlandson, R. F., Kristy, K., Geer, D. M., and Nichols, A. Patient and staff acceptance of robotic technology in occupational therapy: A pilot study. *Journal of Rehabilitation Research and Development* 28, 2 (1991), 33–44.
8. Freedman, R., Jung, H., and Zilberstein, S. Plan and activity recognition from a topic modeling perspective. In *Proceedings of the 24th International Conference on Automated Planning and Scheduling* (2014).
9. Fugl-Meyer, A., Jääskö, L., Leyman, I., Olsson, S., and Steglind, S. The post-stroke hemiplegic patient: a method for evaluation of physical performance. *Scandinavian journal of rehabilitation medicine* 7, 1 (1975), 13–31.
10. Griffiths, T. L., Steyvers, M., Blei, D. M., and Tenenbaum, J. B. Integrating topics and syntax. In *Proceedings of the 18th Annual Conference on Neural Information Processing Systems* (2004).
11. Hogan, N., Krebs, H. I., Charnnarong, J., Srikrishna, P., and Sharon, A. MIT-MANUS: A workstation for manual therapy and training I. In *Proceedings of the IEEE International Workshop on Robot and Human Communication* (1992).
12. Huynh, T., Fritz, M., and Schiele, B. Discovery of activity patterns using topic models. In *Proceedings of the 10th International Conference on Ubiquitous Computing* (2008).
13. Jung, H., Baird, J., Choe, Y., and Grupen, R. A. 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).
14. Jung, H., Baird, J., Choe, Y., and Grupen, R. A. 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).
15. Jung, H., Choe, Y., and Grupen, R. A. Extended virtual presence of therapists through home service robots. In *Proceedings of the 10th ACM/IEEE International Conference on Human-Robot Interaction* (2015).
16. Jung, H., Takahashi, T., Choe, Y., Baird, J., Foster, T., and Grupen, R. A. Towards extended virtual presence of the therapist in stroke rehabilitation. In *Proceedings of the IEEE International Conference on Rehabilitation Robotics* (2013).
17. Jung, H., and Yoo, D. Towards ethical research practice: Anticipating social consequences of robot-mediated rehabilitation. In *Proceedings of the IEEE International Symposium on Ethics in Engineering, Science, and Technology* (2014).
18. Koppula, H. S., and Saxena, A. Learning spatio-temporal structure from rgb-d videos for human activity detection and anticipation. In *Proceedings of the International Conference on Machine Learning* (2013).
19. Kwakkel, G., Kollen, B., and Krebs, H. Effects of robot-assisted therapy on upper limb recovery after stroke: a systematic review. *Neurorehabilitation and Neural Repair* 22, 2 (2008), 111–121.
20. Perry, J. C., Balasubramanian, S., de Pablo, C. R., and Keller, T. Improving the match between ability and challenge: toward a framework for automatic level adaptation in game-based assessment and training. In *Proceedings of the IEEE International Conference on Rehabilitation Robotics* (2013).
21. Prange, G., Jannink, M., Groothuis-Oudshoorn, C., Hermens, H., and IJzerman, M. Systematic review of the effect of robot-aided therapy on recovery of the hemiparetic arm after stroke. *Journal of Rehabilitation Research and Development* 43, 2 (2006), 171–184.
22. Shirzad, N., and der Loos, M. V. Adaptation of task difficulty in rehabilitation exercises based on the user’s motor performance and physiological responses. In *Proceedings of the IEEE International Conference on Rehabilitation Robotics* (2013).
23. Wallach, H. M. Topic modeling: Beyond bag-of-words. In *Proceedings of the 23rd International Conference on Machine Learning* (2006).
24. Wang, G., Zhang, Y., and Li, F.-F. Using dependent regions for object categorization in a generative framework. In *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition* (2006).
25. Wang, Y., and Mori, G. Human action recognition by semi-latent topic models. *IEEE Transactions on Pattern Analysis and Machine Intelligence Special Issue on Probabilistic Graphical Models in Computer Vision* 31, 10 (2009), 1762–1774.
26. Zarit, S. H., and Zarit, J. M. *Mental Disorders in Older Adults: Fundamentals of Assessment and Treatment*. The Guilford Press, New York, NY, 2011.
27. Zhang, H., and Parker, L. 4-dimensional local spatio-temporal features for human activity recognition. In *Proceedings of the International Conference on Intelligent Robots and Systems* (2011).