LEARNING FROM BIOFEEDBACK
Patient-specific Games for Neuromuscular Rehabilitation

Ouriel Barzilay and Alon Wolf
Biorobotics and biomechanics Laboratory, Technion - Israel Institute of Technology, Haifa, Israel

Keywords: Artificial neural networks, Patient-specific rehabilitation, Virtual reality, Biofeedback.

Abstract: Rehabilitation tasks are generally subjected to the physiotherapist's qualitative interpretation of the patient's pathology and needs. Motivated by the recently increasing use of virtual reality in rehabilitation, we propose a novel approach for the design of those biomechanical tasks for an improved patient-specific and entertaining rehabilitation. During training, the subject wears 3D goggles in which virtual tasks are displayed to him. His kinematics and muscles activation are tracked in real time and an inverse model is estimated by artificial neural networks. The resulting inverse model produces a physical exercise according to the observed abilities of the subject and to the expected performance dictated by the physiotherapist. The system offers several advantages to both the patient and the physiotherapist: the tasks can be presented in the form of interactive personalized 3D games with augmented feedback, stimulating the patient's motivation and reducing the need of constant monitoring from the therapist. Additionally, offline quantitative data from every training session can be stored for further analysis. The results of our study on arm movements suggest an improvement in the training efficiency by 10% for the biceps and by 32% ($p=0.02$) for the triceps.

1 INTRODUCTION

Physiotherapy aims at helping patients recover maximal movement and functionality after surgical operations, injuries or strokes. Neuromuscular rehabilitation is generally performed in the form of biomechanical tasks, designed to restore a cognitive or mechanical function within the patient. These tasks are elaborated by the physiotherapist, based on his diagnosis of the patient's pathology. Success of the rehabilitation training relies on the adequate design of these tasks, on the repetition of the physical exercises by the patient, on the subject's motivation and on the feedback to the patient (Holden, 2005). Furthermore, for good results, the task must be adapted to the actual performance of the patient. This adaptive physiotherapy is very difficult to perform online and is subjected to the trainer's interpretation of the patient's performance.

Today, rehabilitation in virtual reality has become a large field of research and several studies have been published on the efficiency and the advantages offered by this approach. Virtual rehabilitation consists in the execution of biomechanical tasks in virtual environments, generally by the means of display devices, biofeedback or haptic instrumentation and adapted software. Virtual rehabilitation has proved efficient in the treatment of neurological diseases (e.g. in Jack et al., 2001 or Holden et al., 2005) for patients with balance disorders (Jacobson et al., 2001) or sports medicine. Studies have shown scientific evidence that motor skills can be learned in virtual environments (Regian et al., 1992) and transferred to the real world (Holden and Dyar, 2002). Furthermore, the augmented feedback on performance offered by virtual reality improves the results of rehabilitation (Shea and Wulf, 1999). It is also likely to increase the motivation of the patient during training (Maclean et al., 2000, Rizzo and Kim, 2005). Some researchers even claim that motor learning in virtual environments can surpass training in the real world (e.g. Todorov et al., 1997).

Training in virtual environments also permits to enhance the rehabilitation platform with computational models and learning systems. In this paper, we introduce a virtual adaptive biofeedback rehabilitation approach aimed at improving neuromuscular training using artificial neural networks able to learn from biofeedback and to

Barzilay O. and Wolf A..
LEARNING FROM BIOFEEDBACK - Patient-specific Games for Neuromuscular Rehabilitation.
DOI: 10.5220/0003679801680174
In *Proceedings of the International Conference on Neural Computation Theory and Applications* (NCTA-2011), pages 168-174
ISBN: 978-989-8425-84-3

produce online new patient-specific virtual physiotherapy missions. With the help of a motion capture system and electromyograms (EMG), our rehabilitation system tracks at any time the kinematics of a subject and his muscle activation. The subject is exposed, via a head mounted display unit, to virtual tasks, which he is asked to perform. After a calibration phase, a neural network is trained to respond to the subject's biofeedback information, based on the desired muscle activation and motions prescribed by the physiotherapist. Once training is complete, the network calculates a new trajectory as biomechanical exercise, subjected to the previous performance of the subject. This adaptive loop is repeated continuously, resulting in an online biofeedback-based adaptive virtual rehabilitation system.

We exploit in this study the ability of artificial neural networks to accurately model systems on which little information is available or complex systems where a computational model is preferred over an explicit model based on arbitrary assumptions and constraints. Biological systems undoubtedly are the most difficult systems to model, as they often involve several functional elements. For instance, virtual upper-limb rehabilitation not only involves motor learning and control of the subject, but also his interpretation of the virtual environment or his hand-eye coordination faculty. Our research hypothesis is that an artificial neural network can be utilized to model biological systems by defining only the input and output signals to the system. This contrasts with works aiming at building sophisticated internal models of human motor control (e.g. Kawato 1990, 1999). We validate our hypothesis on a simple case of neuromuscular rehabilitation. To the best of our knowledge, the approach presented in this paper is novel and has no antecedent in the literature. The results of this study can encourage others in this field to further explore the ability of learning systems to model human functions by the means of biofeedback instrumentation.

2 METHODS

2.1 Experimental Setup

Using the Vicon™ motion capture system capabilities, we track the subject's motions in real time and continuously gather kinematic data. Markers are placed on the subject's body or on part of it (we first focused our study on the arms). During training, the subject is immersed in a virtual environment in which he is shown floating targets (Figure 1). The subject is then asked, for instance, to follow the motions of a target with his pointing finger (*virtual ball* application). His motions are continuously recorded by the system while following the virtual missions presented to him. Moreover in order to record the subject muscular activation, we place electromyograms sensors on key muscles associated related to the motion.

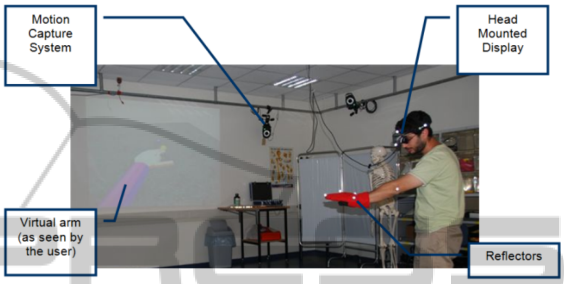

Figure 1: Experimental Setup.

The user receives, in real time, augmented feedback on his performance during training. For instance in the *virtual ball* application, in which the subject must continuously get his hand as near as possible to the ball in motion, the feedback on the distance to the target is provided in several manners. On the virtual replica of the subject's hand is drawn a ball, having the same radius as the virtual ball to reach. The color of this ball changes with respect to the distance to the target, according to a pre-defined color code. The subject can also read his score rising proportionally to the distance to the target. An additional indication on the distance to the virtual ball in the horizontal plane is provided by the shadows of both balls, projected vertically onto the virtual ground. Those indicators allow the user to correct his gestures while performing the exercise. For an improved accuracy yet, when the subject is relatively close to the target, an additional feedback is given to the subject, as the volume of a musical background varies according to the distance to the target. This last estimator offers the subject the opportunity to perform fine-tuning on his hand's position. The activation of a specific muscle is also displayed as sweat drops coming out the virtual sleeve, proportionally to the produced effort.

2.2 An Inverse Model of the Subject

The subject can be seen as a model receiving a biomechanical exercise as an input and producing a performance, that may include kinematic signals (e.g. the trajectory of a limb), as well as EMG signals as an output. We wish to develop a system able to generate a patient-specific physiotherapeutic task, given the kinematic and/or muscular performance of the subject. This goal may be attained from an estimation of the inverse model of the subject, with the desired performance as the input and the exercise trajectory at the output (Figure 2).

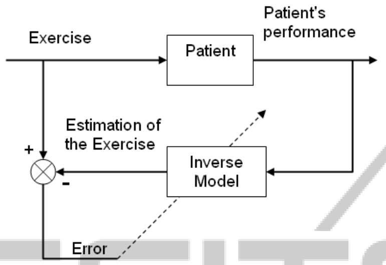

Figure 2: Best Estimated Inverse of the Subject.

2.3 The Learning System

Our goal is to train an artificial neural network capable of producing a subject-specific biomechanical task, given a desired subject's performance. The definition of the subject's performance is arbitrary and may include for example kinematic, kinetic, or muscular parameters.

2.3.1 Network with Kinematic Input Only

In the first phase of our study, we tracked the kinematics of the subject's pointing finger. We use this data to train a neural network. However, we first have to determine the network architecture.

The *universal approximation theorem for neural networks* states that every continuous function that maps intervals of real numbers to an output interval of real numbers can be approximated at any level of desired accuracy by a multi-layer feed-forward neural network with a single hidden layer having a sigmoid activation function. In our case, the network is designed to map, at each instant, the desired position and velocity of a marker placed on the subject, to another spatial point and velocity corresponding to the displayed exercise. Consequently, we need a system capable of modeling the mapping from R^6 to R^6. We assume the mapping to be a continuous function and use the *approximation theorem* to build a multi-layer feed-forward network with one hidden layer. This network has a generic architecture for all subjects and is specific to this definition of the performance. Nonetheless, each subject will have his own tuned network.

To define the network's architecture, we start with a known exercise trajectory. The subject is presented with this task and is asked to follow the target trajectory displayed to him. Concurrently, the tracking system records his motions at given timestamps. The recorded kinematic data serves as an input set to the network, while the corresponding exercises displayed to him serve as target outputs. We use Levenberg-Marquardt error back-propagation learning method (Moré, 1977) so that the error between the actual output of the network and the target output is minimized. The network's weights are initialized according to Nguyen-Widrow's method (Nguyen & Widrow, 1990).

The network used for a single marker contains six input neurons and six output neurons: corresponding to the spatial position and velocity vectors. Next, the number of neurons in the hidden layer needs to be determined. On the one hand, this parameter affects the runtime and should therefore be minimized. On the other hand, it also influences the system's accuracy and balance between precision and efficiency must be attained. We performed a set of tuning experiments where several healthy subjects were given a cyclic trajectory as a task to follow while their motions were tracked (Figure 3).

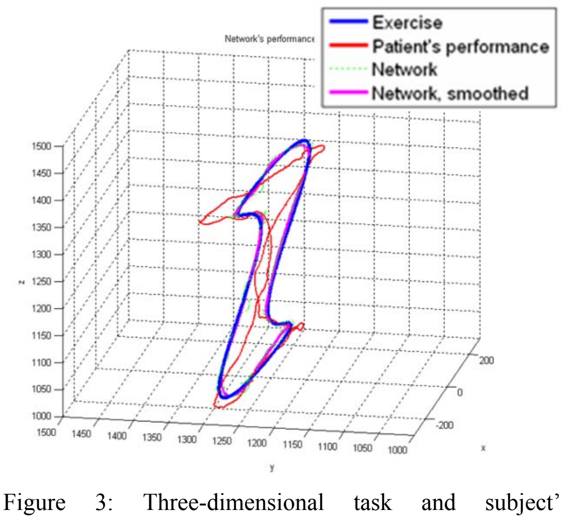

Figure 3: Three-dimensional task and subject's performance.

The network was trained on the training set composed by the exercise and the average performance over the number of cycles. We used three criteria to evaluate the network's performance for each subject and exercise: the output error, the positional output error and the average of the positional output errors over each of the ten different cycles performed by the subject. The average of the output errors reflected the network's ability to extrapolate its results on samples that were not directly included in the training set. After iteratively testing different sigmoid activation functions for the hidden layer of the network, numbers of epochs and numbers of hidden neurons, the final configuration of the network was determined.

The resulting network comprised seven hidden neurons with the hyperbolic tangent as activation function and the number of epochs was set to 50. The reader is referred to previous publication (Barzilay and Wolf, 2009) for a detailed explanation on the setting of the network's architecture.

2.3.2 Network with Kinematic and EMG signals as Input

In the second step of this research, we added the patient's biceps and triceps EMG signals to the input of the network, such that the new exercise would be designed with respect to the information on the muscles of the subject as well as the knowledge on the kinematic data of his limb.

The data provided by the electromyograms contain useful information that can be deciphered by signal processing. There are numerous ways described in the literature to extract this information from the EMG signal, including analysis in time or frequency domains. We use the workflow described in Hodges and Bui (1996) to compute the linear envelope of the signal by processing it in the time domain. The processed signal is needed at every instant in our application and the processing time has to be minimized, all the more since several signals are needed simultaneously. We accelerated this operation by using the processed signals from the precedent instant and reduced the processing time by approximately 96% (Barzilay and Wolf, 2011). This fast implementation allows providing the subject with continuous visual feedback on his own muscular performance during the training.

The same parameters that were described in section 2.3.1 are used to evaluate the network, but now in addition to the 3D curve, the desired EMG performance specific to that trajectory should also be designed. We therefore determined a few desired cyclic trajectories for the limb of the subject and recorded the EMG performances of a dozen of healthy subjects. The average of this set of data is then used as the desired EMG performance over a specific trajectory, and fed as input to the neural network together with the trajectory of the desired kinematic performance.

The number of neurons in the hidden layer has been set to 17, according to the evaluation criteria which were previously used.

2.3.3 System Evaluation

The first network, described in section 2.3.1, considers only the endpoint kinematics of the subject and has obviously less physiotherapeutic interest than the network involving the subject's muscular performance (section 2.3.2). Nevertheless, the optimistic results (section 3.1 and Barzilay and Wolf, 2009) suggested evidence of the feasibility of modeling human motor control with neural networks and brought us to expand the subject model to include muscular performance as well.

From a therapeutic perspective, the muscles activation of the patient is more significant than his ability to accurately reproduce specific trajectories. For that reason, we focus our efforts on minimizing the error in the EMG performance, whereas the kinematics error is considered more moderately. Although the EMG signals are calibrated from measurement of the maximal voluntary contraction prior to the training, the signals' amplitudes tend to differ between different subjects. Furthermore, we focus on the rhythmical patterns of the muscles more than on the activation intensity. To do that, we consider the error between the desired and actual EMG performance in the frequency domain.

For the evaluation of the adaptive system, we thus consider the root mean squared deviation, in the frequency domain, between the desired EMG performance and the smoothed EMG performance of the subject. Each participant (n = 16) performed motor training on two exercises: the patient-specific exercise produced by the trained neural network (*adapted training*), and a general exercise having for trajectory the desired kinematic performance. The latter resembles a standard physiotherapeutic session where the physiotherapist demonstrates to the patient, for example with his hand, the desired gesture to reproduce (*conventional training*). The primary criterion for the system evaluation was defined as the ratio of the errors obtained in the performances in both cases.

3 RESULTS

3.1 Network with Kinematic Input

The exercise trajectory, designed by the network, deviates by 15 millimeters per point in average from the exact trajectory. This average deviation is reduced to 3-5 millimeters per point when a smoothing filter is applied to the trajectory produced by the network. The network succeeds by such to estimate the inverse model of the subject.

In Figure 3, one can see that, due to the relative location between the planar target and the subject's eye, the task (in blue) can be perceived as a projection on a plane normal to the subject's line of sight. Nevertheless, the neural network system learned and corrected the projection, although far from being a linear phenomenon.

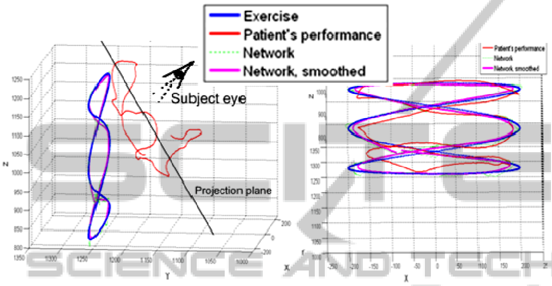

Figure 4: The exercise projection: side (left) and front (right) views.

Once trained, the network is capable of generating a patient-specific exercise, given a desired patient performance. The inverse model of the patient can be evaluated by comparing the measured performance of the subject with respect to desired one. Given in figure 5 are the desired performance, the patient-specific exercise created by the network, and the performance of the subject on this adapted task.

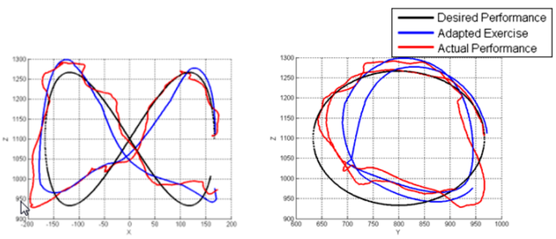

Figure 5: Desired Vs. actual performance: front (left) and side (right) views.

Let us recall that, during rehabilitation session, the subject does not see the whole exercise trajectory, but only the virtual ball, in motion along that trajectory. Moreover, he is not exposed to the desired trajectory. In the depicted case, the average distance between the subject's performed trajectory and the desired performance was approximately equal to the average distance between the hand trajectory and the trajectory of the displayed exercise. However, in several sections of the task, the subject's trajectory was noticeably closer to the desired trajectory than to the virtual ball, as can be observed on the left side of the side view in Figure 5. It is also notable that, in this section, the network was able to predict that, in order to cause the subject's hand to follow the desired trajectory (in black), the virtual ball had to be displayed a bit farther along the *y* axis (farther from the subject). This observation suggests that the model created by the network was able to detect some of the subject's behavioral patterns. This phenomenon was observed in several sessions and for different subjects. It is also notable, in Figure 5, that the system's prediction was effected twice on the same portion of the exercise trajectory, while the subject's hand had different velocities.

3.2 Network with Kinematic and EMG Signals as Input

Figure 6 demonstrates the capability of the network to adapt itself to the muscular information recorded from the electromyograms. Patterns characteristic to the desired signal appear within the performance of the subject on the exercise designed by the system. Frequency-domain analysis shows how close the spectra of the desired and the actual performance signals are. We found that in many cases the subject obtained a better performance on the network-designed exercise than if he is directly shown the desired trajectory of his limb as an exercise (Figure 6).

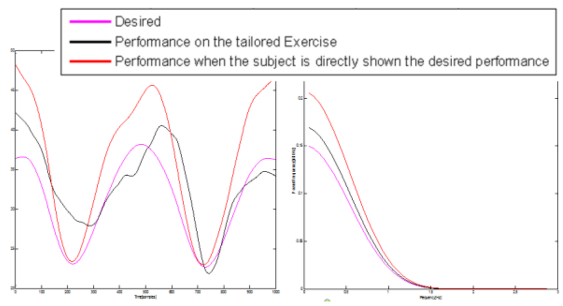

Figure 6: EMG performance: time (left) and frequency (right) domains.

The deviations from the desired EMG performance in the frequency domain for the adapted and conventional training are presented in Table 1. These results are summarized in Table 2.

Table 1: Adapted and conventional training comparison.

Suhject	Adapted Training		Conventional Training		Error Ratio	
	Biceps	Triceps	Bi.	Tri.	Bi.	Tri.
#1	53.13	66.56	73.84	77.13	1.39	1.16
#2	36.25	85.10	76.95	79.76	2.12	0.94
#3	59.49	32.24	78.75	58.90	1.32	1.83
#4	85.00	85.64	40.44	92.74	0.48	1.08
#5	88.92	60.19	67.09	83.63	0.75	1.39
#6	63.85	97.61	56.58	84.20	0.89	0.86
#7	92.61	94.98	62.03	89.72	0.67	0.94
#8	62.94	39.51	18.53	68.19	0.29	1.73
#9	72.05	59.18	81.32	88.28	1.13	1.49
#10	74.67	42.19	148.20	58.98	1.98	1.40
#11	72.05	59.18	81.32	88.28	1.13	1.49
#12	65.00	29.28	74.27	86.06	1.14	2.94
#13	92.78	64.54	103.48	48.57	1.12	0.75
#14	56.12	55.39	71.95	69.94	1.28	1.26
#15	73.25	47.07	54.98	52.16	0.75	1.11
#16	44.71	93.98	54.85	72.04	1.23	0.77

Table 2: Adapted and conventional training comparison - Summary.

	Adapted Training		Conventional Training		Error Ratio	
	Biceps	Triceps	Bi.	Tri.	Bi.	Tri.
Average	68.30	63.29	71.54	74.91	**1.10**	**1.32**
Std. Dev.	16.50	22.53	28.10	14.16	0.48	0.54

Most subjects (n = 14, 87.5 %) benefitted from our system for at least one of the muscles, and almost half of them improved the accuracy of their muscular performance for both biceps and triceps (n = 7, 43.75%). In summary, the results indicate that the average muscular performance of the subjects is closer to the desired performance when the exercise is generated by the system, rather than set as the desired kinematic performance like in conventional physiotherapy. This is indicated by a 10% increase for the biceps performance and by 32% for the triceps performance.

A one-tailed Student T-test shows that the improvement in the triceps performance is attained with statistical significance ($p = 0.02$). However, the improvement in the biceps performance is lesser in magnitude and in statistical significance ($p = 0.34$ and $p = 0.09$ with omission of two subjects). We believe that this is due to the fact that the physical exercise stimulated by the system involves the biceps in a smaller measure than the triceps or the shoulder muscles.

4 DISCUSSION

We introduce, in this study, a platform for motor and cognitive rehabilitation, able to model the subject's kinematics and to generate a subject-specific physiotherapeutic exercise. The system requires no prior knowledge on the patient, nor any model of his motor control or trajectory planning. It only involves the desired performance dictated by the physiotherapist and a training reference set, recorded in situ from the patient's performance prior to rehabilitation. To date, and to the best of our knowledge, no study combining virtual reality rehabilitation and learning algorithms for patient-specific training has been reported.

The developed system offers several opportunities to both the physiotherapist and the patient. The virtual tasks can be designed as interactive games and stimulate the motivation of the patient during rehabilitation. We have developed several applications where the subjects are enjoined to pop bubbles, stop soccer balls, or whack objects with their hands in a controlled way. Most of the participants expressed their enthusiasm after having performed motor training in our virtual applications.

In every session, all the kinematic and EMG data are stored and may be further analyzed offline by the physiotherapist. Furthermore, the system proved to emphasize some kinematic and muscular patterns in motor training, and may contribute to a better diagnosis of the subject. This system may find its application in patients after stroke, with cerebral palsy, dyslexia or other developmental coordination disorders.

At this time, we have tested the system on healthy subjects only, choosing to generalize and validate it before testing it on pathological subjects. The success of the system in learning some of the subject’s behavioral patterns leads us to expect good results in the modeling of motor patterns in patients with pronounced pathology. Besides clinical trials, we would like to expand the system to combine more markers and EMG sensors, and other biometric sensors.

Besides physiotherapy, this system could prove useful in sportive performance enhancement, in the development of new types of human-machine interfaces, in entertainment, and in the training of any kind of motor skills.

While a description of the *subject model*, including his motor control and trajectory planning, hand-eye coordination and probably many additional features, would be very difficult to elaborate, the computational power of the simplest form of feed-forward neural networks provided very optimistic results in the modeling of the subject. First to combine virtual rehabilitation with machine learning of human models, the positive results of this study encourage carrying on the use of biofeedback-based artificial intelligence and virtual reality, for applications in therapy and other diverse areas.

REFERENCES

Barzilay, O. and Wolf, A. (2009). An adaptive virtual system for neuromuscular rehabilitation. In: IFMBE Proc. World Congress on Medical Physics and Biomedical Engineering, Munich, Germany. 25(4): 1291–94.

Barzilay, O. and Wolf, A. (2011). A fast implementation for EMG signal linear envelope computation. *Journal of Electromyography and Kinesiology*, 21: 678-682.

Hodges, P. and Bui B. (1996). A comparison of computer-based methods for the determination of onset of muscle contraction using electromyography. *Electroenc. Clin. Neurophysiol,* 101, 511–519.

Holden, M. K. (2005). Virtual environments for motor rehabilitation: review. *Cyberpsychology & Behavior,* 8(3),187–211.

Holden, M. K. and Dyar, T. (2002). Virtual environment training: a new tool for neurorehabilitation. *Neurology Report,* 26, 62–71.

Holden, M. K., Dyar, T., Schwamm L. and Bizzi, E. (2005). Virtual Environment-Based Telerehabilitation in Patients with Stroke. *Presence: Teleoperators and Virtual Environments*, 14(2), April, 214-233.

Jack, D., Boian R., Merians, A. S., Tremaine, M., Burdea, G. C., Adamovich, S. V., Recce, M. and Poizner, H. (2001). Virtual reality–enhanced stroke rehabilitation. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 9, 308–318.

Jacobson, J., Redfern, M. S., Whitney, S. L. and Wilson, J. B. (2001). Balance NAVE: A Virtual Reality Facility for Research and Rehabilitation of Balance Disorders. In *Proc. Virtual Reality Software and Technology Meeting*, Banff, Canada.

Kawato, M. (1990). Computational schemes and neural network models for formulation and control of multijoint arm trajectories. In *Neural networks for control*, eds. Miller WT, Sutton RS and Werbos PJ. MIT press.

Kawato, M. (1999). Internal models for motor control and trajectory planning. *Current Opinion in Neurobiology* 9:718–727.

Maclean, N., Pound, P., Wolfe, C. and Rudd, A. (2000). A qualitative analysis of stroke patients' motivation for rehabilitation. *BMJ,* 321, 1051–1054.

Moré, J. J. 1977. The Levenberg-Marquardt algorithm: Implementation and Theory. In: *Numerical Analysis. Lecture Notes in Mathematics 630*, ed. Watson GA, 105–116. Berlin: Springer-Verlag.

Nguyen, D., Widrow, B. (1990). Improving the learning speed of 2-layer neural networks by choosing initial values of the adaptive weights. In *Proc.International Joint Conference on Neural Networks*, July, 3:21–26.

Regian, J. W., Shebilske, W. L. and Monk, J. M. (1992). Virtual reality: an instructional medium for visual spatial tasks. *Journal of Communication,* 7, 131–145.

Rizzo, A. and Kim, G. J. (2005). A SWOT Analysis of the Field of Virtual-Reality Rehabilitation and Therapy. *Presence: Teleoperators and Virtual Environments*, 14(2), April, 119-146.

Shea, C. H. and Wulf, G. (1999). Enhancing motor learning through external-focus instructions and feedback. *Human Movement Science,* 18, 553–571.

Todorov, E., Shadmer, R. and Bizzi, E. (1997). Augmented Feedback Presented in a Virtual Environment Accelerates Learning of a Difficult Motor Task. *Journal of Motor Behavior*, 29, 147–158.

VICON Motion Systems at http://www.vicon.com/