Proceedings of the 28th IEEE
EMBS Annual International Conference
New York City, USA, Aug 30-Sept 3, 2006

Artificial Neural Network Prediction Using Accelerometers to Control Upper Limb FES During Reaching and Grasping Following Stroke

Phil Tresadern, Sibylle Thies, Laurence Kenney, David Howard
Centre for Rehabilitation and Human Performance Research,
University of Salford, Salford, UK, M6 6PU
Email: p.tresadern@salford.ac.uk

John Y. Goulermas
Dept. of Electrical Engineering and Electronics,
University of Liverpool, UK, L69 3GJ
Email: j.y.goulermas@liverpool.ac.uk

Abstract— **This work investigates arm acceleration as a control signal for Functional Electrical Stimulation (FES) of the upper limb during reaching and grasping. We segment the reach and grasp motion into phases and present an Artificial Neural Network (ANN) approach that estimates the phase of the reaching cycle from accelerometer signals. We then select the stimulator command that maximizes successful triggering without unnecessary risk to the patient's safety. Our results suggest that the algorithm successfully generalizes between sessions and patients but is less successful at generalizing between different motions.**

I. Introduction

Following upper motor neuron damage, motor function can be restored through the use of functional electrical stimulation (FES). However, the control of such systems remains a major challenge. Systems controlled by sources such as Electromyogram (EMG) and Electroencephalogram (EEG) signals have been demonstrated although such signals are limited in terms of transduction, which is sensitive to external noise, and current surface EEG approaches remain limited by information transmission rates [1].

As an alternative to electrophysiological signal-based approaches, a number of groups have used body segment motion to control both FES [2] and powered prosthetic limbs [3]. Contra-lateral shoulder motion as a control signal has proven clinically successful in the Freehand system, used following high level spinal cord injury [2], but is less than ideal for subjects with a lower level of impairment.

The motion of a segment on the same limb has been used to control foot drop stimulation [4] but applying a similar approach to the upper limb is more challenging. Despite the existence of characteristic movement patterns, kinematic redundancy of the upper limb and the under-constrained nature of reaching make it difficult to distinguish such motions from other types of movement. Further complications arise due to increased variability in reaching kinematics following upper motor neuron damage and variation in the desired endpoint of the reach. The use of an artificial motion to distinguish control gestures from other movement has been shown to be feasible, if cosmetically less than ideal [5].

As a first approximation, the input signals must be processed in such a way as to generate a binary (on/off) output. In gait analysis, where phase transitions are well-defined and motion is highly periodic, Machine Learning (ML) has proven successful for several years. Popular tools include Artificial Neural Networks [6], ANFIS [6], [7], Inductive Logic [6], [4] and Adaptive Logic Networks [4]. Only in recent years, however, have similar methods been applied to the upper limb where motion is less constrained. Even then, results have only been demonstrated using electrophysiological measurements [8] and joint angular velocities measured using flexible goniometers [9].

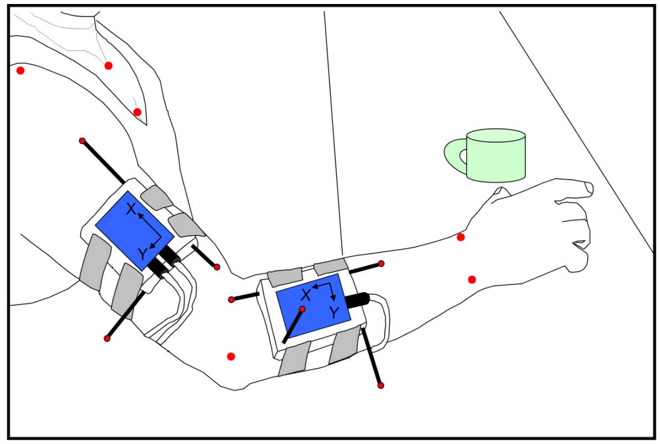

Fig. 1. Experimental set up: the arm is adorned with markers (for rigid body motion capture) and Inertial Measurement Units (for validation).

We propose an approach, developed as part of the EU-funded Healthy Aims project (www.healthyaims.org), where we segment the "reach and grasp" motion into three phases – Reach/Retract (RR), Grasp/Release (GR) and Manipulate (M) – and use an ANN to estimate probabilities corresponding to our belief that the observed motion "belongs to" each phase. However, in many cases it is desirable to "err on the side of caution" when deciding whether to stimulate the hand *e.g.* stimulating whilst holding a hot cup of coffee presents a threat to the safety of the user. Therefore, we minimize the *risk* associated with each possible decision (neglected in previous upper limb studies) using a *loss matrix*.

To our knowledge, this work is one of the first to use Machine Learning with accelerometer signals to control upper-limb FES, based on natural (as opposed to artificial [5]) movement. In contrast to other methods that require the entire motion history before prediction can take place [9], our system is currently being prepared for real time trials on stroke patients.

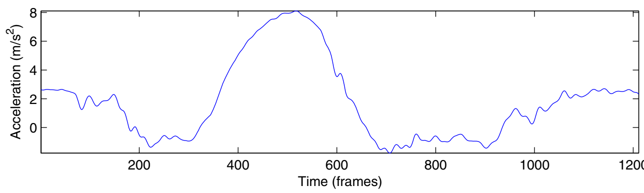

Fig. 2. Synthesized acceleration along the axis of the arm.

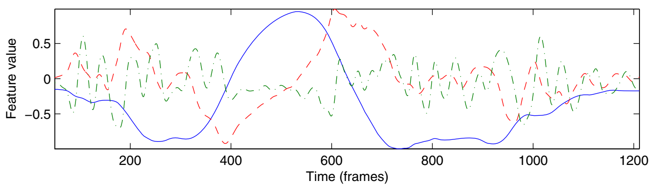

Fig. 3. First three normalized features extracted from original signal. Note that no features are available until at least one complete 'window' (50 measurements, in this case) have been observed.

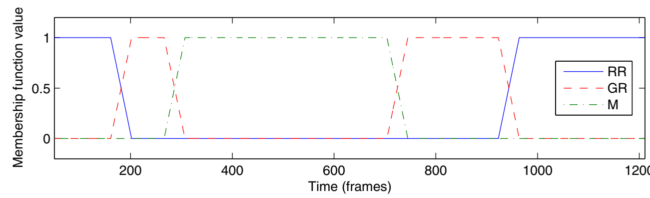

Fig. 4. Assigned "fuzzy" phase labels (membership function values).

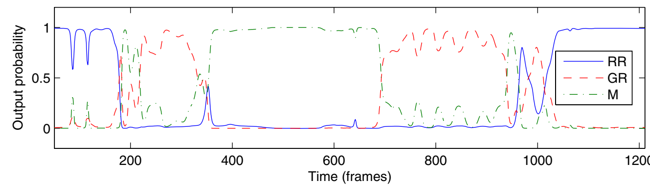

Fig. 5. Example output probabilities produced by the neural network. This example corresponds to an average loss of 0.097 and the resulting stimulation has a mean squared error of 0.125.

II. Data Collection

Ethical approval for the study was obtained and two hemiplegic stroke patients recruited. Both subjects were more than six months post-stroke, medically stable and right-handed with right upper limb function compromised. Both were able to reach forward towards objects at a comfortable distance without significant discomfort or pain and had a modified Ashworth scale of 3 or below in the flexors of the affected hand, wrist and elbow. Neither subject had fixed contractures of the elbow, wrist and fingers. Patient 1 was a 32 year old female and patient 2 was an 83 year old male.

Movement data of the right upper limb were collected using a 10 camera motion capture system (Vicon, Oxford Metrics, UK). Rigid clusters of markers were placed on the torso, upper arm and forearm (see Fig. 1), and additional markers were placed at anatomical locations (*e.g.* humeral epicondyles). The measured marker positions were then used to infer joint centres and rigid body motion of the limbs.

In order to evaluate system performance independently of errors caused by sensor misalignment during donning and doffing, we synthesize accelerometer signals [10] from the computed rigid body motions at a sampling rate of 100Hz. These synthesized signals, validated against actual measurements recorded from an inertial unit (XSENS MT-9), were then used as inputs to the algorithm.

III. Neural Network Design

The Artificial Neural Network was implemented using the Neural Network Toolbox in Matlab (The Mathworks, Inc) and had an input layer of 4 nodes, one hidden layer of 5 *tansig* nodes and an output layer of 3 *softmax* nodes. For computational reasons, the network was trained using Scaled Conjugate Gradients (*trainscg*) with 80% of the data used for training and 20% for validation to avoid overfitting.

For these experiments, we used only the acceleration along the axis of the forearm (Fig. 2). Although it was possible to use this raw signal as input to the ANN, we used a window of 50 samples (corresponding to 0.5s) to exploit the time history of the signal. In order to reduce the dimensionality of this (largely redundant) data, we computed a 4D feature vector using coefficients from the Discrete Cosine Transform (DCT) of the original signal. For numerical stability, the coefficients were then normalized to lie in the range $[-1, 1]$ using the *premnmx* function (see Fig. 3) and used as inputs to the ANN.

Although studies have been conducted into hand opening during reaching tasks for healthy adults [11], the timings of events for post-stroke reaching are not well defined. Therefore, we identified phase transitions by inspection, assigning a "fuzzy" label to each instant in time (see Fig. 4) to reflect uncertainty close to phase transitions. These fuzzy labels then served as target outputs from the ANN.

A. Risk-based decision making

Using the *softmax* transfer function at the output layer ensures that the output of the network can be interpreted as a vector of probabilities (see Fig. 5), each element reflecting our confidence that the observed motion belongs to the corresponding phase. However, simply selecting a stimulation level based on which phase has the highest probability does not take into account the risk associated with each decision. For example, if "Reach/Grasp" and "Manipulate" were assigned equal probability, we must select the phase that would present the smallest risk (*i.e.* Manipulate) should our decision be incorrect.

We designed the system according to three specifications:

1) For comfort and to minimize power requirements, we *prefer* stimulation to be *off* during "Reach/Retract"
2) For functionality, we *require* stimulation to be *on* during "Grasp/Release"
3) For safety, we *demand* that stimulation be *off* during "Manipulate"

To implement these specifications, we defined a loss matrix such that each type of misclassification incurred a different penalty according to the risk presented to the patient's safety. The penalty for classifying the phase correctly is zero

TABLE I

MISCLASSIFICATION PENALTIES. WE DENOTE THE TRUE PHASE BY X AND THE ESTIMATED PHASE BY Y.

		X: RR	X: GR	X: M
	RR	None	Medium	Low
Y	GR	Medium	None	High
	M	Low	Medium	None

whilst a small penalty is incurred if the phase is misclassified with no difference in stimulation (*e.g*. Reach/Retract as Manipulation). A medium penalty is incurred if misclassification results in stimulation but with no risk to patient safety (*e.g*. Reach/Retract as Grasp/Release) whereas a large penalty is incurred for any misclassification that may endanger the patient (*e.g*. Manipulation as Grasp/Release). Table I expresses these rules in tabular form where we imposed low, medium and high penalties of 0.25, 0.5 and 1.0, respectively.

Given the three probabilities estimated by the ANN, we computed the expected loss associated with each possible decision and selected the phase that incurred the smallest penalty. This admitted a margin of error such that we could "err on the side of caution" in critical regions. Stimulation was 'on' only for the Grasp/Release phase of the motion and 'off' at all other times.

IV. EXPERIMENTAL METHOD

In order to evaluate the system, we quantified generalization ability with respect to test session, subject and motion. These cases represent small, medium and large variations in the presented inputs. We used several training data sets to quantify not only generalization ability but also how performance was affected by training on data from a mixture of sessions compared with a single session.[1]

We used motion analysis data from 5 test sessions (A-E), as described in Table II, each containing 7 repetitions of the described motion. The motions were designed according to an experimental protocol such that the position of the patient and the endpoint of the reach were fixed between test sessions for repeatability analysis. In each experiment the neural network was trained on several datasets, each consisting of six motions. Each training set contained either six motions from one test session, three motions each from two test sessions, or two motions each from three test sessions. Each test set consisted of one motion from each relevant session.

To quantify accuracy and generalization ability, we employed two error metrics: the mean squared error ('MSE') directly compares algorithm output with the desired stimulation and therefore does not take into account the different penalties applied to various misclassifications; the average loss ('Loss') over the test motion addresses this shortcoming by penalizing *all* misclassifications according to the loss matrix.

[1]This has practical importance in order to determine whether training a bespoke network presents a significant benefit over a "one-size-fits-all" system that can be used in off-the-shelf operation.

TABLE II

DATA SETS.

Dataset	Patient	Day	Motion
A	1	1	Lift glass to mouth, palm down[a]
B	1	2	Lift glass to mouth, palm down
C	2	3	Lift glass to mouth, palm down
D	1	1	Lift glass off table, palm sideways
E	1	1	Move plate laterally, palm up

[a]Denotes the orientation of the palm at the start of the motion.

Since the neural network was initialized with random weights and the error function typically has many local minima, we trained every network forty times and computed the mean and standard deviation of the error. This averaged out fluctuations between tests as a result of random initialization.

A. Generalization over session, subject and motion

In the first of three experiments, we trained on data sets A and B, corresponding to the same subject performing the same motion on different days, in order to identify whether the algorithm is sensitive to small changes in motion as a result of day-to-day variation (due to fatigue, spasticity *etc*.).

A second test investigated the robustness of the algorithm to medium variations in the motion patterns. We trained on data sets A and C, taken from different subjects performing the same motion. Due to the relative rigidity of the experimental protocol, the motion patterns are similar with some variation due to the "style" of the subject.

Finally, we investigated whether the neural network could generalize to novel movements. This challenging case was trained on data sets A, D and E, providing signals from the same subject performing different motions. Note that the motions in data sets A and D are relatively similar (lifting a glass into the air from a tabletop) whereas the motion in data set E is very different (moving a plate laterally).

V. RESULTS

A. Generalization with respect to session

Table III shows that there was little difference in MSE between the two test days. This is not unexpected since the data is from the same subject performing the same motion. Of course, this error is likely to vary inversely with respect to how repeatably the subject can perform a given motion.

TABLE III

GENERALIZATION OVER TEST SESSION.

Training Set		Test Motion A	Test Motion B
A(6[a])	MSE	0.129 (0.020)[b]	0.129 (0.013)
	Loss	**0.100** (0.018)	0.104 (0.010)
B(6)	MSE	0.127 (0.087)	**0.117** (0.083)
	Loss	0.113 (0.066)	**0.098** (0.061)
A(3),B(3)	MSE	0.125 (0.017)	**0.120** (0.016)
	Loss	0.105 (0.021)	**0.102** (0.016)

[a]Number of motions used from dataset

[b]Mean (Std Dev)

B. Generalization with respect to subject

Compared to the previous results, Table IV shows that the algorithm did not generalize as well between patients due to the increased variation in motion. However, we note that errors were generally higher for motion C, suggesting that repeatability is lower for this patient and motion.

TABLE IV
GENERALIZATION OVER SUBJECT.

		Test Motion	
Training Set		A	C
A(6)	MSE	**0.129** (0.020)	0.193 (0.016)
	Loss	**0.100** (0.018)	0.184 (0.014)
C(6)	MSE	**0.111** (0.021)	0.165 (0.047)
	Loss	0.151 (0.017)	**0.139** (0.022)
A(3),C(3)	MSE	**0.109** (0.009)	0.172 (0.007)
	Loss	**0.101** (0.015)	0.166 (0.011)

C. Generalization with respect to motion

From the results shown in Table V, we see that the neural network generalized relatively well between the similar motions (data sets A and D) but did not generalize well for the dissimilar motion (data set E). This is indicated by the large errors present in the final column, with the exception of when the network was trained on data solely from that motion. Furthermore, as the training set became more general the average error increased for all motions, suggesting that a different neural network may be required for highly dissimilar motions.

TABLE V
GENERALIZATION OVER MOTION.

			Test Motion	
Training Set		A	D	E
A(6)	MSE	**0.129** (0.020)	0.174 (0.040)	0.274 (0.033)
	Loss	**0.100** (0.018)	0.149 (0.025)	0.246 (0.024)
D(6)	MSE	**0.123** (0.041)	0.187 (0.035)	0.282 (0.032)
	Loss	**0.127** (0.038)	0.143 (0.031)	0.289 (0.024)
E(6)	MSE	0.255 (0.012)	0.366 (0.021)	**0.164** (0.023)
	Loss	0.230 (0.005)	0.265 (0.009)	**0.093** (0.020)
A(3),D(3)	MSE	**0.112** (0.032)	0.121 (0.052)	0.262 (0.023)
	Loss	**0.093** (0.025)	0.122 (0.032)	0.243 (0.022)
A(3),E(3)	MSE	**0.188** (0.047)	0.261 (0.056)	0.247 (0.051)
	Loss	**0.135** (0.026)	0.204 (0.035)	0.166 (0.047)
D(3),E(3)	MSE	**0.182** (0.124)	0.240 (0.097)	0.250 (0.116)
	Loss	0.195 (0.080)	0.217 (0.057)	**0.161** (0.079)
A(2),D(2),E(2)	MSE	**0.132** (0.041)	0.165 (0.066)	0.261 (0.036)
	Loss	**0.108** (0.031)	0.154 (0.042)	0.203 (0.045)

D. Discussion

We make a number of observations from these results: the 'MSE' and 'Loss' error metrics sometimes disagree over which which motion was best recognized although the 'Loss' metric tends to agree more closely with our intuition; motion A was apparently the easiest to recognize, as evidenced by consistently lower errors in most cases (where data from motion A was included during training); performance typically deteriorated when trained on multiple subjects although in some cases error *decreased* when trained on mixed data, possibly as a result of improved generalization.

VI. Conclusion

We have presented an application of Machine Learning for upper limb FES triggering, based on forearm acceleration. Stimulation was selected in such a way as to minimize risk and the correct stimulation was observed for 80-90% of the motion cycle. A neural network, trained on various datasets, generalized well over test sessions and (to a lesser extent) patients. Generalization over motion was demonstrated for similar motions and less so for dissimilar ones.

Potential developments of the system include: incorporating temporal smoothness using tools such as a Finite State Machine; employing the loss matrix during training to increase accuracy in critical regions; optimizing network structure using genetic algorithms [12]; investigating other Machine Learning methods (*e.g.* non-parametric methods); employing alternative features (*e.g.* wavelet coefficients).

VII. Acknowledgments

The authors gratefully acknowledge the assistance of Christine Smith and Julie Rigby during data collection.

References

[1] T. M. Vaughan, W. J. Heetderks, L. J. Trejo, W. Z. Rymer, M. Weinrich, M. M. Moore, A. Kubler, B. H. Dobkin, N. Birbaumer, E. Donchin, E. W. Wolpaw, and J. R. Wolpaw, "Brain-computer interface technology: a review of the second International Meeting," *IEEE Trans. on Neural Systems and Rehabilitation Engineering*, vol. 11, pp. 94–109, 2003.

[2] J. Hobby, P. N. Taylor, and J. Esnouf, "Restoration of tetraplegic hand function by use of the neurocontrol Freehand system," *Journal of Hand Surgery (British and European Volume)*, vol. 26, pp. 459–464, 2001.

[3] J. A. Doubler and D. S. Childress, "An analysis of extended physiological proprioception as a prosthesis-control technique," *Journal of Rehabilitation Research and Development*, vol. 21, pp. 5–18, 1984.

[4] R. Williamson and B. J. Andrews, "Gait event detection for FES using accelerometers and supervised machine learning," *IEEE Trans. on Rehabilitation Engineering*, 2000.

[5] K. Y. Tong, A. F. T. Mak, and W. Y. Ip, "Command control for functional electrical stimulation hand grasp systems using miniature accelerometers and gyroscopes," *Medical and Biological Engineering and Computing*, vol. 41, no. 6, pp. 710–717, Nov. 2003.

[6] S. Jonic, T. Jankovic, V. Gajic, and D. Popovic, "Three machine learning techniques for automatic determination of rules to control locomotion," *IEEE Trans. on Biomedical Engineering*, vol. 46, no. 3, pp. 300–310, Mar. 1999.

[7] R. T. Lauer, B. T. Smith, and R. R. Betz, "Application of a neuro-fuzzy network for gait event detection using electromyography in the child with cerebral palsy," *IEEE Trans. on Biomedical Engineering*, vol. 52, no. 9, pp. 1532–1543, Sept. 2005.

[8] J. P. Giuffrida and P. E. Crago, "Functional restoration of elbow extension after spinal-cord injury using a neural network-based synergistic FES controller," *IEEE Trans. on Neural Systems and Rehabilitation Engineering*, vol. 13, no. 2, pp. 147–152, June 2005.

[9] L. Cenciotti, S. Micera, M. C. Carrozza, P. Dario, and M. Popovic, "A hybrid system for the prediction of upper arm articular synergies using statistical and soft-computing techniques," in *Proc. Int'l Funtional Electrical Stimulation Society*, 2001.

[10] S. B. Thies, P. Tresadern, L. Kenney, D. Howard, Y. Goulermas, and C. Smith, "A "virtual sensor" tool to simulate accelerometer output for upper limb FES triggering," in *Proc. World Conference of Biomechanics*, 2006.

[11] R. G. Marteniuk, J. L. Leavitt, C. L. MacKenzie, and S. Athenes, "Functional relationships between grasp and transport components in a prehension task," *Human Movement Science*, vol. 9, no. 2, pp. 149–176, Apr. 1990.

[12] J. Y. Goulermas, A. H. Findlow, C. J. Nester, D. Howard, and P. Bowker, "Automated design of robust discriminant analysis classifier for foot pressure lesions using kinematic data," *IEEE Trans. on Biomedical Engineering*, vol. 52, no. 9, pp. 1549–1562, Sept. 2005.