

Reconstructing muscle activation during normal walking: a comparison of symbolic and connectionist machine learning techniques

Ben W. Heller^{1,*}, Peter H. Veltink², Nico J. M. Rijkhoff³, Wim L. C. Rutten², Brian J. Andrews^{1,**}

¹ Bioengineering Unit, University of Strathclyde, Glasgow, Scotland

² Biomedical Engineering Division, Department of Electrical Engineering, University of Twente, Enschede, The Netherlands

³ Department of Urology, Faculty of Medicine, Katholic University of Nijmegen, Nijmegen, The Netherlands

Received: 15 July 1992/Accepted in revised form: 30 March 1993

Abstract. One symbolic (rule-based inductive learning) and one connectionist (neural network) machine learning technique were used to reconstruct muscle activation patterns from kinematic data measured during normal human walking at several speeds. The activation patterns (or desired outputs) consisted of surface electromyographic (EMG) signals from the semitendinosus and vastus medialis muscles. The inputs consisted of flexion and extension angles measured at the hip and knee of the ipsilateral leg, their first and second derivatives, and bilateral foot contact information. The training set consisted of data from six trials, at two different speeds. The testing set consisted of data from two additional trials (one at each speed), which were not in the training set. It was possible to reconstruct the muscular activation at both speeds using both techniques. Timing of the reconstructed signals was accurate. The integrated value of the activation bursts was less accurate. The neural network gave a continuous output, whereas the rule-based inductive learning rule tree gave a quantised activation level. The advantage of rule-based inductive learning was that the rules used were both explicit and comprehensible, whilst the rules used by the neural network were implicit within its structure and not easily comprehended. The neural network was able to reconstruct the activation patterns of both muscles from one network, whereas two separate rule sets were needed for the rule-based technique. It is concluded that machine learning techniques, in comparison to explicit inverse muscular skeletal models, show good promise in modelling nearly cyclic movements such as locomotion at varying walking speeds. However, they do not provide insight into the

biomechanics of the system, because they are not based on the biomechanical structure of the system.

1 Introduction

Biomechanical models have been described for normal walking (Crowninshield and Brand 1981; Patriarco et al. 1981; Hatze 1984; Davy and Audu 1987), walking of paraplegics assisted by functional electrical stimulation (FES) of paralysed muscles (Yamaguchi and Zajac 1990), and for prosthetic gait (Koopman 1989). These models are based on explicit models of the musculoskeletal system that describe the relationships between the neural activation of muscles and the kinematics of movement, using the theories of multi-body dynamical systems and detailed knowledge of the anatomy and contractile properties of muscles. However, despite the sophistication of such models, they are restricted in their ability to reconstruct accurately individual muscle forces and activation patterns due to the many assumptions inherent in them, for instance, that a movement will be performed so as to minimise certain metabolic energy or muscle fatigue criteria. Furthermore, it appears to be difficult to identify and track the time-varying system parameters.

It can be argued that the most optimal model of a human movement exists within the neurological control mechanisms of the person performing the movement; for instance, most able-bodied individuals constantly consciously and subconsciously hone and adapt their gait to produce optimal walking patterns. The measure of optimality appears to be a dynamic combination of factors including efficiency, speed, safety, normality, minimisation of pain, etc. The motor control sequences (or rules) are a combination of spinal reflexes, patterns stored in the lower brain centres and cerebral cortex (Grillner 1975; Bussel et al. 1988) and conscious intentions. They are thus impossible to articulate fully – a problem which is acknowledged by psychologists studying motor skills (Anderson 1980; Davids and Myers

* *Present address:* Medical Physics, Royal Hallamshire Hospital, Sheffield, England

** *Present address:* Glenrose Rehabilitation Hospital, Edmonton, Alberta, Canada

Correspondence to: P. H. Veltink, Biomedical Engineering Division, Department of Electrical Engineering, University of Twente, P.O. Box 217, 7500 AE Enschede, The Netherlands

1990). This is a situation analogous to the so-called Feigenbaum bottleneck of artificial intelligence, meaning the difficulty of manually eliciting knowledge from an expert. The solution adopted in the artificial intelligence domain is to use an alternative way of obtaining the knowledge, namely learning by example, in which the expert system produces its own rules from general principles extracted from the training data presented to it by the expert. This inductive approach may produce more effective rules than those explicitly articulated by the expert (Michalski and Chilauski 1980).

Given a set of data containing examples, each consisting of a vector of input values (the attributes) and a desired output value vector (which may consist of continuous values, or of discrete classes), inductive learning techniques will attempt to find relationships between the inputs and the outputs. In this way, the system can be said to learn from the data presented to it. The system's memory consists of either decision rules or, for a neural network, synaptic weights.

We will compare two kinds of inductive learning algorithms in this report. The first kind is called symbolic or rule-based inductive learning. In this technique the program attempts to produce (or learn) a set of if-then-else decision rules which may be presented as a decision tree – a directed graph representing the sequentially ordered decision rules produced by the program.

The second technique is the use of a multi-layer feedforward perceptron network with supervised learning-via-error back-propagation (Rumelhart and McClelland 1986; Lippmann 1987). Because of the supervised learning, this technique may also be considered an inductive learning method.

The performance of the two machine-learning methods will vary with the type of data they are tested on because their function is determined by their structure. Neural networks are better at X of N functions (where at least X features of a possible N must be present for an example to be a member of a class); symbolic learning techniques require very complex rule trees to solve these problems. Symbolic methods are better for singly sufficient attribute values of the "if a, then b" kind (Mooney 1990).

A thorough comparison of the performance of symbolic learning (ID3) and connectionist (error back-propagating neural network algorithm) techniques on real-world data was performed by Shavlik et al. (1991). They used five data sets to assess the algorithms' performances, reaching the following conclusions:

1. The neural network gave slightly higher classification accuracy than ID3 when tested on new examples.
2. The neural network could be significantly better on numerical data sets [this contradicts the findings of Weiss and Kapouleas (1989)].
3. The neural network was slightly more accurate when the examples were noisy or incompletely specified [however, the chi-squared pre-pruning method used for ID3 was probably less effective than post-pruning techniques (Mingers 1989)].

Other differences were that ID3 could be trained 150 times faster, and it classified new examples 10 times faster than the neural network [although the error back-propagation technique may not be the fastest learning mechanism (Willis et al. 1990)]; the neural-network architecture (number of hidden layers, number of neurons) needed to be selected, which is considered more of an art than a science, and perhaps most importantly for the present application, the comprehensibility of the knowledge acquired by the symbolic learning technique was much higher, because the rules which are obtained can be interpreted.

In this study a symbolic or rule-based inductive learning method is compared with a neural network method for their ability to reconstruct muscle activation patterns from kinematic data at several walking speeds. The motivation for this work is twofold:

1. To compare the use of rule-based inductive learning techniques with neural networks for the induction of rules in the general domain of movement analysis.
2. To investigate the possibility of cloning the rules used by a human expert during normal gait at varying speeds. It provides a method of directly synthesising control strategies for robot and FES-induced gait, without having to find an inverse model of the system (the model can be said to reside within the rules or weights generated by the inductive learning strategy).

This study was limited to reconstructing the activation patterns (the classes) of two muscles (semitendinosus and vastus medialis) from the bilateral foot contact signals and the joint angles of the ipsilateral hip and knee and their first and second derivatives with respect to time (the attributes). The measurements were made on one unimpaired subject walking at two different speeds.

2 Methods

2.1 Experimental methods

Experimental data were obtained from one male, unimpaired individual by instrumenting hip and knee unilaterally with flexible goniometers (Penny & Giles Ltd). These angular data were sampled at 50 Hz by an Analog Devices RTI-815 12-bit resolution data acquisition card installed in an IBM-PC compatible computer. They were then low-pass filtered at a cut-off frequency of 5 Hz by a finite impulse response (FIR) filter, then doubly differentiated using a second FIR filter. Foot-floor contact was determined by monitoring electrical contact between metal plates under the feet and a conducting rubber floor; this signal was then integrated to give information about the time spent in that phase of the gait cycle (see Fig. 1).

Activation patterns of the ipsilateral semitendinosus muscle (one of the hamstring muscles, responsible for knee flexion) and the vastus medialis muscle (a knee extensor) were determined by means of bipolar surface EMG measurements. The signals were rectified and low-pass filtered using a second-order Butterworth filter with

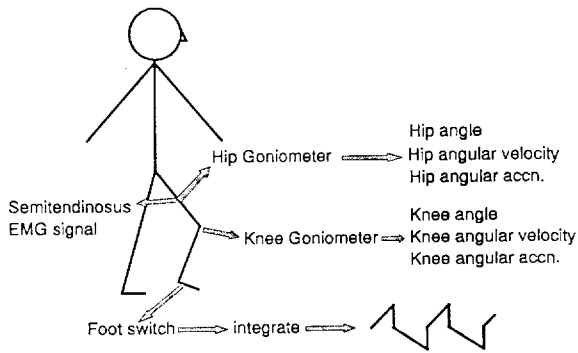


Fig. 1. Kinematic and muscle activation signals during reciprocal gait were measured from a healthy man. Hip angle and knee angle were measured using flexible goniometers. From these signals the angular velocities and angular accelerations (accn.) were estimated. Foot contact was measured with foot switches. This binary signal was integrated to give information about the time spent in that phase of the gait cycle. Muscle activation of semitendinosus and vastus medialis were estimated from rectified and low-pass filtered electromyography

a 25-Hz cut-off frequency, prior to being sampled at 200 Hz. The EMG signals were further low-pass filtered off-line, using a 57 coefficient FIR filter with a 5-Hz cut-off frequency and zero phase shift; this gave an envelope with a bandwidth similar to that measured from the kinematic joint angle data. The subject performed four trials of four or five double steps at each of two speeds: natural speed and a faster speed; this was to evaluate the ability of the machine learning techniques to generate the correct activation patterns at varying walking speeds. The data were recorded once a steady-state gait had been established. At each speed the filtered data (hip and knee angles, angular velocities, angular accelerations and integrated foot contacts) from three trials were used as a training set for both inductive learning methods. The data from the remaining trials at each speed were used as the testing set for analysing the performance of both methods in reconstructing the activation patterns from the kinematic data.

2.2 Symbolic learning

The algorithm used is based on the hierarchical mutual information classifier of Sethi and Sarvarayudu (1982). The program that implements this algorithm, *Empiric*, is based on the *Disciple* program described by Kirkwood et al. (1989), but has been rewritten to include fuzzy class membership (described below) and to make more efficient use of the limited memory available in 640k personal computer. It can accommodate up to 10^4 examples, each having a maximum of sixteen, 16-bit attributes.

The muscle activation signal was divided into five levels (an arbitrary choice), and each of these was taken as an independent class. The assignment of the continuous data into these arbitrary classes imposes unjustified divisions on the data; e.g. a muscle activation slightly to one side of a quantisation level is considered to be in a different class to one that is just beyond the level. This may prevent the algorithm from choosing a classification

threshold that combines both examples into the same class, even though the threshold may be justified from a consideration of the continuous data. We thus chose to assign class membership on a fuzzy basis (Zadeh 1965): each example is considered to be both a member of the class it is in and of the next nearest class. The fuzzy membership value depends on the position of the example within the class and is given by:

$$\mu(x, u) = 1 - \frac{|x - M(u)|}{Q} \quad \text{for } |x - M(u)| < Q$$

$$\mu(x, u) = 0 \quad \text{otherwise} \quad (1)$$

where $\mu(x, u)$ is the fuzzy membership value of an example of value x in class u , Q is the size of a quantum (the difference between adjacent quantisation levels), and $M(u)$ is the value of the midpoint of class u .

$M(u)$ is given by:

$$M(u) = (u - 1)Q. \quad (2)$$

The algorithm performs a hierarchical partitioning of the attribute space. Each new node of the decision tree contains a rule based on a threshold of one of the attributes. With each new rule the example set is further subdivided. The decision tree will continue to grow until each leaf (terminal node) contains members of only one class. For non-exact, non-mutually exclusive or noisy data this will lead to a very large decision tree containing non-robust rules which reflect the noise in the data rather than genuine causal relationships. To avoid this overfitting, the algorithm can be made to terminate the formation of new nodes before all the examples have been correctly classified, thus limiting the size of the tree. The choice of tree size depends on the characteristics of the data being classified, and is in general determined by trial and error. For a fuller discussion of this see Clark (1990).

Two different rule sets were induced from the training data: one which had 3 rules (7 nodes) and a more complex one which had 31 rules (63 nodes). Both rule sets were then used to classify each example in the testing sets, and the reconstructed and actual classes were compared.

2.3 Neural network

Neural network computations were performed on an IBM-PC compatible computer using the N Works software package (Neural Works Inc.). A multi-layer perceptron network with supervised learning via back-propagation was used in this study (Lippmann 1987). The inputs to the neural network consisted of the same attributes and trials as for the symbolic inductive learning training set detailed above. The time order of the examples was randomised before they were applied to the network, and the attribute values were normalised between -1 and 1 . The desired outputs were the continuous muscle activation levels rather than the quantised levels used above. At the beginning of the training session the weights of the nodes were initialised randomly between -0.1 and 0.1 . The input signals and the corresponding output activation patterns were applied to the network at each time step, and the weights of the nodes were adapted via the back-propagation training

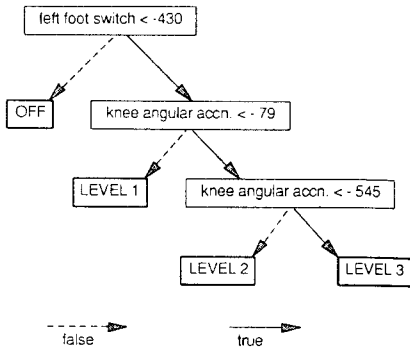


Fig. 2. Small decision tree for rule-based inductive learning method to reconstruct semitendinosus activation during normal walking. Only the left foot switch data were used for on-off timing of the activation burst and knee angular acceleration (accn.) for determining the level of activation in the burst

algorithm. The learning coefficient was set to 0.9 and the smoothing factor to 0.6 (Rumelhart and McClelland 1986). The complete training set was applied to the network 30 times. Two different networks were considered, one with one hidden layer and one with two hidden layers.

3 Results

As an example of the output rules that were generated, the rule set for the (simpler) 7-node network reconstructing semitendinosus is shown in Fig. 2. This small decision tree only uses the information of the left foot switch for timing of the activation burst and the knee angular acceleration for determining the activation level during the burst.

The synthesised muscle activation levels using symbolic inductive learning are plotted together with the actual levels for both walking speeds in Figs. 3a, 4a, 5a







Fig. 3a-d. Measured and reconstructed activation patterns of semitendinosus for slow walking. **a** applying a small fuzzy rule base (3 rules; 7 nodes); **b** applying a large fuzzy rule base (31 rules; 63 nodes); **c** applying a small neural network (one hidden layer); **d** applying a large neural network (two hidden layers)

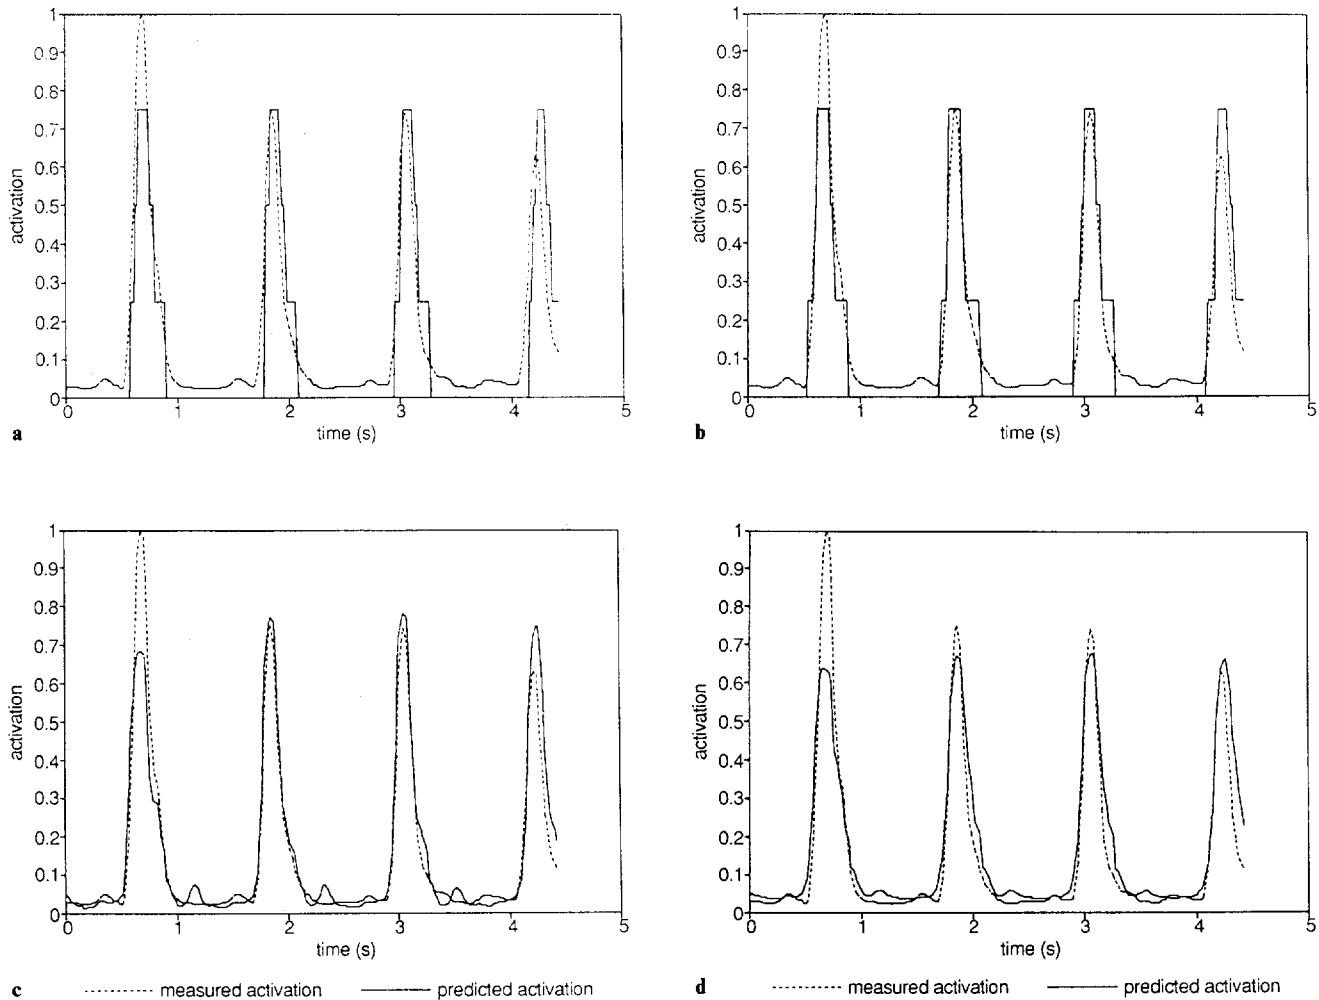


Fig. 4a–d. Measured and reconstructed activation patterns of semitendinosus for fast walking. **a** applying a small fuzzy rule base (3 rules; 7 nodes); **b** applying a large fuzzy rule base (31 rules; 63 nodes); **c** applying a small neural network (one hidden layer); **d** applying a large neural network (two hidden layers)

and 6a for the small rule set and Figs. 3b, 4b, 5b and 6b for the larger rule set. The equivalent outputs from the neural networks for both speeds are given in Figs. 3c, 4c, 5c and 6c for the smaller network and Figs. 3d, 4d, 5d and 6d for the larger network.

To assess the ability of the techniques to reconstruct the activation patterns, the following comparisons between reconstructed and measured outputs were made for all techniques:

1. The difference in the onset and offset times of measured and reconstructed activation bursts are given in Table 1. The times at which the activation crossed 20% of the maximum value of each measured burst were determined. These quantities indicate the accuracy of the timing of the reconstructed activity bursts.
2. The average difference in the time integral of each reconstructed and measured activation burst between the onset and offset times described above are given in Table 2. This quantity indicates the accuracy of the total

activity of the reconstructed bursts, which is related to the total associated mechanical impulse.

These aspects (accuracy of timing and integrated value of the muscle activity) determine the accuracy of the associated mechanical output of the musculoskeletal system, because the movement of the slow and mainly inertial skeletal system is mainly determined by the timing and size of the mechanical impulse which is generated by the relatively short muscle-burst activity. Therefore, timing and integral value of this activity is more important than the actual time course of the activity.

All the techniques had similar timing performances. No large systematic errors were found in the onset and offset timings, since the averages and the standard deviations of the timing errors were of the same order of magnitude as the standard deviations of the timing errors. The average and standard deviations of the timing errors were worse for vastus medialis than for semitendinosus in the case of slow walking. This can be

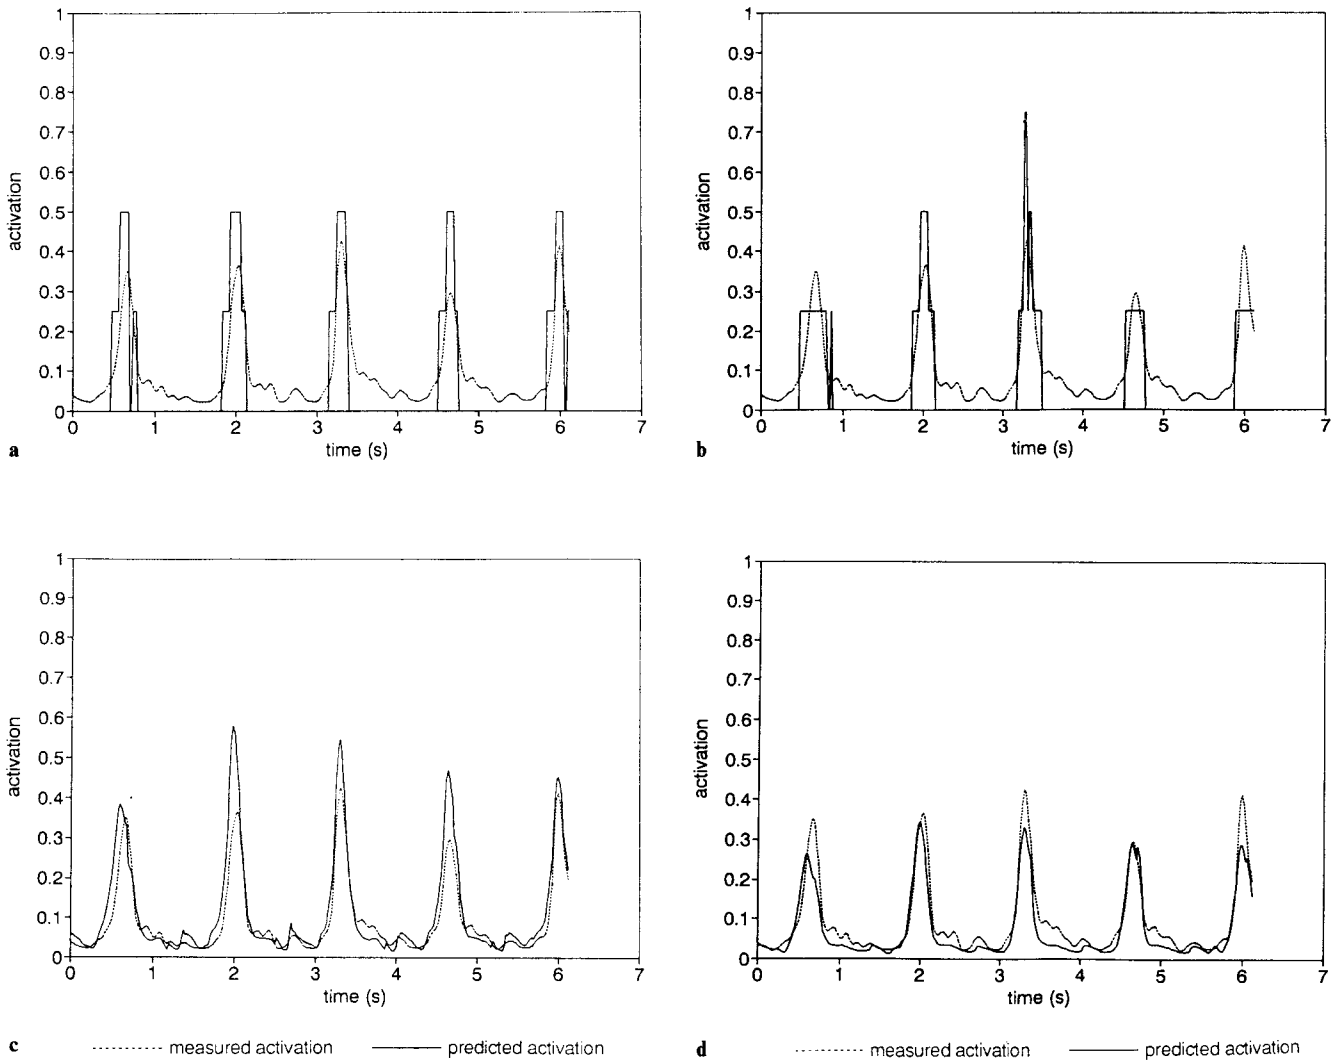


Fig. 5a–d. Measured and reconstructed activation patterns of vastus medialis for slow walking. **a** applying a small fuzzy rule base (3 rules; 7 nodes); **b** applying a large fuzzy rule base (31 rules; 63 nodes); **c** applying a small neural network (one hidden layer); **d** applying a large neural network (two hidden layers)

explained by observing that vastus muscle does not turn on and off as sharply as semitendinosus, and thus a small difference in the amplitude threshold at which the reconstruction turns on will lead to a larger difference in timing.

The errors in the reconstruction of the integrated activities were similar for all four techniques (Table 2). Also, no clear differences between the errors for both walking speeds and muscles can be distinguished.

Figures 3–6 show some variations between subsequent reconstructed activation bursts of the same trial. These variations did not always follow correctly the variations between the corresponding measured activation bursts.

4 Discussion

The results (Figs. 3–6, Tables 1, 2) indicate that both inductive learning methods were able to distinguish

between the two walking speeds: the activation patterns for slow and fast walking were reconstructed equally well in terms of onset times, offset times and integrated burst activities.

In order to judge the acceptability of the timing errors, they should be related to the dynamical characteristics of the musculoskeletal system: the dynamics of muscles can be characterised by a rise time of approximately 50–100 ms, depending on the muscle type. The resonance frequency of the skeletal system depends on the phase of the walking cycle (stance or swing), but can be estimated to be lower than 1 Hz. Therefore, timing errors on the order of 50 ms seem to be acceptable.

The errors in the integrated activation values appear to be rather large, because the relative errors in the integrated activation (up to an average of 32%) are directly related to an error in the corresponding mechanical impulse delivered by the muscle. One possible reason




Fig. 6a–d. Measured and reconstructed activation patterns of vastus medialis for fast walking. **a** applying a small fuzzy rule base (3 rules; 7 nodes); **b** applying a large fuzzy rule base (31 rules; 63 nodes); **c** applying a small neural network (one hidden layer); **d** applying a large neural network (two hidden layers)

for these relatively large errors might be that the system is insufficiently observable with the limited sensor set used: the differences between the cycles might not be sufficiently represented in the measured kinematic data. For example, movements other than in the sagittal plane are not measured.

A possible explanation for the observation that the timing of the activation patterns appeared to be reconstructed with sufficient accuracy, but not the size (integral of the burst), is that more information is needed for a good reconstruction of the size than for a good reconstruction of the timing. This is in concert with the rules of the small rule set depicted in Fig. 2: the rule set only with one rule relied on the timing information contained within the integrated foot-contact signal to determine if the muscle was on or off; the other two rules were needed to determine the magnitude of the activation from the angular acceleration at

the knee. It is also consistent with the results of Arendt-Nielsen et al. (1991), who found that the shapes (timings) of the EMG activation and knee kinematic patterns (normalised to the gait cycle period) remained similar for walking at different speeds, whereas the amplitudes varied.

Whereas a deductive model, explicitly describing the inverse musculoskeletal dynamics, will be valid (within the limits of its assumptions) for all movement patterns, the induced models will only be valid over a range of movements that are similar to ones contained within their training sets. Also, the induced models do not give biomechanical insight into the system, because they only relate input and output signals and are not based on the biomechanical structure of the system like that of the deductive biomechanical models.

However, the accuracy of the timing and size of the reconstructed muscle activity bursts do not appear to be

Table 1. Errors in the onset and offset timing of the reconstructed EMG bursts compared with original bursts (average \pm standard deviation)

Technique	Speed	Semitendinosus		Vastus medialis	
		Onset (ms)	Offset (ms)	Onset (ms)	Offset (ms)
Small rule set	Slow	-16 ± 26 (n = 5)	8 ± 23 (n = 5)	0 ± 51 (n = 5)	-190 ± 96 (n = 4)
Large rule set	Slow	-16 ± 26 (n = 5)	8 ± 23 (n = 5)	32 ± 39 (n = 5)	-135 ± 93 (n = 4)
Small neural network	Slow	12 ± 23 (n = 5)	8 ± 23 (n = 5)	-68 ± 23 (n = 5)	-85 ± 72 (n = 4)
Large neural network	Slow	-24 ± 22 (n = 5)	8 ± 23 (n = 5)	4 ± 33 (n = 5)	-145 ± 87 (n = 4)
Small rule set	Fast	20 ± 29 (n = 4)	40 ± 35 (n = 3)	53 ± 23 (n = 3)	-7 ± 31 (n = 3)
Large rule set	Fast	-30 ± 12 (n = 4)	40 ± 35 (n = 3)	7 ± 12 (n = 3)	13 ± 42 (n = 3)
Small neural network	Fast	-10 ± 12 (n = 4)	27 ± 50 (n = 3)	0 ± 0 (n = 3)	0 ± 20 (n = 3)
Large neural network	Fast	-20 ± 0 (n = 4)	33 ± 42 (n = 3)	20 ± 0 (n = 3)	-20 ± 20 (n = 3)

The activation threshold applied for determination of the onset and offset moments is 20% of the maximal value of the measured activation bursts. A negative error implies that the reconstructed burst is ahead of the measured burst

Table 2. Errors in the integrated values of the reconstructed activation bursts in comparison with the measured activation bursts (average \pm standard deviation)

Technique	Speed	Semitendinosus Integrated activation (%)	Vastus medialis integrated activation (%)
Small rule set	Slow	32 ± 17 (n = 5)	14 ± 20 (n = 4)
Large rule set	Slow	14 ± 17 (n = 5)	3 ± 21 (n = 4)
Small neural network	Slow	7 ± 14 (n = 5)	30 ± 20 (n = 4)
Large neural network	Slow	27 ± 13 (n = 5)	-22 ± 11 (n = 4)
Small rule set	Fast	6 ± 21 (n = 3)	-30 ± 13 (n = 3)
Large rule set	Fast	16 ± 23 (n = 3)	14 ± 28 (n = 3)
Small neural network	Fast	6 ± 23 (n = 3)	21 ± 14 (n = 3)
Large neural network	Fast	8 ± 22 (n = 3)	-26 ± 4 (n = 3)

The bursts were integrated between the onset and offset times considered in Table 1. The errors are expressed as a percentage of the integrated values of the corresponding measured bursts. A positive value implies that the area under the reconstructed burst is larger than the area under the measured burst

better in the deductive biomechanical models (Crowninshield and Brand 1981; Patriarco et al. 1981; Davy and Audu 1987). These biomechanical models are complex and contain many parameters which are difficult to identify. Furthermore, they are based on many assumptions: for instance the models need a criterion to estimate the otherwise undetermined contribution of the individual muscles to the joint torques found by the inverse dynamics (Crowninshield and Brand 1981). Different

criteria and imposed constraints may lead to substantial differences in the reconstructed activation patterns (Patriarco et al. 1981).

An advantage of the use of the inductive techniques presented in this paper over explicit inverse musculo-skeletal models is that they generate the EMG activation level with no time delay. An explicit model's estimation of muscular activation must always lag by a time equal to the delay between electrical activation and mechanical torque generation. As there is no such lag for the inductive learning techniques, they must incorporate anticipatory knowledge of the movements to be performed. This is solely due to the repetitive nature of walking.

Both the rule-based inductive algorithm and the neural-network inductive approach showed comparable performance in the reconstruction of muscle activation patterns from kinematic data. Because of this comparable performance a choice between both methods will depend on other characteristics of the methods.

The advantages of the rule-based inductive learning technique are that the rules are explicit, comprehensible, easily encoded into a knowledge base and may be executed quickly. Furthermore, the algorithm is variably selective: only the "best" (in terms of average mutual information gain) attribute threshold will be selected at any node on the tree; thus, it is possible to identify the most important attributes (sensors) in any rules that are produced.

An advantage of the neural network technique is that one network is able to model the output of more than one muscle. The rule-based inductive learning technique required one independent decision tree per output. Furthermore, whilst the validity of these models cannot be guaranteed for unfamiliar patterns, the sigmoid weighting function used in the neural network ensures that the model will be continuous, i.e. small changes in input will cause small changes in output, thus small deviations in

Machine learning of normal walking patterns

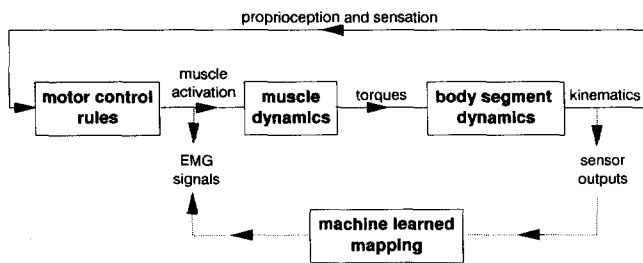


Fig. 7. The mapping of the relationship between the sensor outputs and the EMG signals can either be considered as an *indirect* model of the muscle and body segment dynamics, or a *direct* input-output model of the neural system relating proprioception and sensation to muscle activation

the input pattern should still produce valid outputs. This cannot be said for the rule-based inductive learning technique in which a small change at the inputs can produce a large change at the outputs.

The neural network technique produced continuous outputs, whereas the rule-based inductive learning technique produced discrete, five-level outputs. This is unlikely to be important in practice, as the low-pass filtering characteristics of the muscle and the inertial dynamics of the limb will smooth out the sudden transitions.

The induced rules (or weights for the neural networks) in both inductive methods link the kinematic variables to the muscle activation levels. They can be considered as providing either a direct input-output model (cloning) of the neural system that relates sensory inputs to motor outputs, or an inverse model of the musculoskeletal dynamics (see Fig. 7). Although these two descriptions seem contradictory, they are both equally valid and both equally limited.

Possible applications of the reconstruction of muscular activity from sensor information are in the automatic control of paralysed muscles of paraplegic gait assisted by FES, the manufacturing of low-cost motion analysis equipment that could reconstruct movement patterns from only a few sensors and the control of robot locomotion (Craig 1986). More generally, both techniques are entirely data-driven and contain no causal knowledge of the application domain; thus they are suitable for any classification problem for which a suitable training set can be found.

Acknowledgements. We acknowledge the assistance of Ir. Rob F. M. Kleissen from the Roessingh Rehabilitation Centre in Enschede, The Netherlands, for assisting us in performing the gait analysis measurements using his experimental setup. This work has been partially funded by CAMARC (Computer Aided Motion Analysis in a Rehabilitation Context) project, which is part of the European Community AIM (Advanced Informatics in Medicine) programme. The collaboration which enabled this research has been aided by the Erasmus programme of the European Community.

References

- Anderson JR (1980) Cognitive psychology and its implications. Freeman, San Francisco.
- Arendt-Nielsen L, Sinkjær T, Nielsen J, Kallesoe K (1991) Electromyographic patterns and knee joint kinematics during walking at various speeds. *J Electromyogr Kinesiol* 1:89–95
- Bussel B, Roby-Brami A, Azouvi PH, Biraben A, Yakovlev A, Held JP (1988) Myoclonus in a patient with spinal cord transection: possible involvement of the spinal stepping generator. *Brain* 111:1235–1245
- Clark P (1990) Machine learning: techniques and recent developments. Research memorandum TIRM-90-041, Turing Institute, Glasgow
- Craig JJ (1986) Introduction to robotics. Addison-Wesley, Reading, Mass
- Crowninshield RD, Brand RA (1981) A physiologically based criterion of muscle force prediction in locomotion. *J Biomech* 14:793–801
- Dauids K, Myers C (1990) The role of tacit knowledge in human skill performance. *J Hum Movement Stud* 19:273–288
- Davy DT, Audu ML (1987) A dynamic optimization technique for predicting muscle forces in the swing phase of gait. *J Biomech* 20:187–201
- Grillner S (1975) Locomotion in vertebrates: central mechanisms and reflex intervention. *Physiol Rev* 55:247–304
- Hatze H (1984) Quantitative analysis, synthesis and optimization of human locomotion. *Hum Movement Sci* 3:5–25
- Kirkwood CA, Andrews BJ, Mowforth P (1989) Automatic detection of gait events: a case study using inductive learning techniques. *J Biomed Eng* 11:511–516
- Koopman HFJM (1989) The three-dimensional analysis and prediction of human walking. Ph.D. thesis, University of Twente, The Netherlands
- Lippmann RP (1987) An introduction to computing with neural nets. *IEEE ASSP Mag* 4–22
- Michalski RS, Chilausky RL (1980) Knowledge acquisition by encoding expert rules versus computer induction from examples. *Int J Man Mach Stud* 12:63–87
- Mingers J (1989) An empirical comparison of pruning methods for decision tree induction. *Mach Learning* 4:227–243
- Mooney RJ (1990) Backpropagation versus learning decision trees. *Neural Network Rev* 4:84–85
- Patriarco AG, Mann RW, Simons SR, Mansour JM (1981) An evaluation of the approaches of optimization models in the prediction of muscle forces during human gait. *J Biomech* 14:513–525
- Rumelhart DE, McClelland JL (1986) Explorations in the microstructure of cognition, vol 1. Cambridge, Mass
- Sethi IK, Sarvarayudu GPR (1982) Hierarchical classifier design using mutual information. *IEEE Trans Pattern Anal Mach Intell* 4:441–445
- Shavlik JW, Mooney RJ, Towell G (1991) Symbolic and neural learning algorithms – an experimental comparison. *Mach Learning* 6:111–143
- Weiss SM, Kapouleas I (1989) An empirical comparison of pattern recognition, neural nets and machine learning classification methods. Proceedings of the 11th International Joint Conference on Artificial Intelligence, Detroit, Mich, pp 688–693
- Willis MJ, De Massimo C, Montague GA, Tham MT, Morris AJ (1990) Solving process engineering problems using artificial neural networks. In: McGhee J, Grimble MJ, Mowforth P (eds) Knowledge-based systems for industrial control. (IEE control engineering series 44). Peregrinus, London pp 34–46
- Yamaguchi GT, Zajac FE (1990) Restoring unassisted natural gait to paraplegics via functional neuromuscular stimulation: a computer simulation study. *IEEE Trans Biomed Eng* 37:886–902
- Zadeh LA (1965) Fuzzy sets. *Inf Control* 8:338–353