Multiple Model Adaptive Control of Functional Electrical Stimulation

O. Brend, C. T. Freeman, M. French
University of Southampton, United Kingdom, {owb106/cf/mcf}@soton.ac.uk

***Abstract*—This paper establishes the feasibility of multiple model switched adaptive control within functional electrical stimulation (FES) based upper limb stroke rehabilitation. An estimation-based multiple model switched adaptive control (EMMSAC) framework for nonlinear, time-invariant systems is described, and extensions are presented that enable its application to time-varying Hammerstein structures that can accurately represent the stimulated arm. A principled design procedure is then developed to construct both a suitable set of candidate models from experimental data, and a corresponding set of tracking controllers. The procedure is evaluated by applying it to a sample of able-bodied, young participants in order to produce a general EMMSAC controller. This is then applied to a different sample of the population during an isometric, non-voluntary, trajectory tracking task. Results confirm that multiple model switched adaptive control can provide consistent levels of performance for the population with minimal prior plant model identification. The results show it is possible to eliminate model identification while employing closed-loop controllers that maintain high performance in the presence of rapidly changing system dynamics. This paper hence addresses critical limitations to effective stroke rehabilitation in a clinical setting.**

I. Introduction

Of the fifteen million cases of stroke per annum, approximately a third result in death and another third in permanent disability [1]. In the UK approximately 70% of first time stroke sufferers experience upper limb impairment [2], and less than 40% regain motor function to a level that enables them to perform functional tasks independently [2]. This can limit the ability of sufferers to engage in activities that are frequently performed in day-to-day life (e.g. eating, bathing and dressing). Successful rehabilitation requires repeated practice of tasks to promote motor re-learning. However, conventional rehabilitation, involving manual assistance by a physiotherapist, does not translate to an improved ability to perform functional tasks [3].

Functional electrical stimulation (FES) is a technology that has the potential to provide more efficient, cost-effective therapy, enabling more frequent practice of movements and increased motivation compared to conventional therapy. FES involves the application of an electrical stimulus to impaired muscles to elicit a contraction that produces a desired movement. Its use is supported by clinical studies [4], [5] and neurophysiology and motor learning results [3], [6]. There is evidence that the outcome of FES-based rehabilitation is improved if a patient is attempting to perform a functional task concurrently with the application of electrical stimulation [6]. When FES is applied, electrical impulses travel towards the target muscle, but also towards the spinal cord. Rushton [6] suggests that the coincidental arrival of these impulses, and the voluntary impulses sent from the brain strengthens the synaptic connections within the spinal cord. This synaptic strengthening, termed Hebbian learning, requires maximal voluntary effort from the patient and the coordinated application of electrical stimulation as an aid [6]. Most clinical FES systems are either open-loop or triggered and thus have limited ability to activate impaired muscles in a way that exploits the effect of Hebbian learning. Application of closed-loop control can increase the (tracking) accuracy with which the applied stimulation matches the intended movement via feedback of measured signals such as force, position, velocity and muscle activation (electromyography (EMG)).

For precise closed-loop feedback control of FES, a mathematical description of electrically stimulated muscle is vital. However, identification of such a model is widely seen as impractical in a clinical setting due to time constraints and rapidly changing dynamics (due to fatigue, spasticity and changing physiological and environmental factors such as skin impedance, temperature and electrode placement). One of the few model-based upper limb FES controllers to be employed in a clinical intervention with stroke patients is iterative learning control (ILC) [7]. When used to assist functional reaching tasks, the tracking accuracy provided by ILC gave rise to statistically significant reduction in patient impairment [8], [9] but a lengthy identification procedure was required at the beginning of each treatment session, and unpredictable, rapid variation in the system degraded performance as each treatment session progressed.

The inability to completely adapt to rapidly changing experimental conditions is commonly highlighted as the fundamental stumbling block limiting the transfer of advanced FES control techniques to clinical practice [10]. Poor control accuracy means FES cannot support the intended movement, and the effectiveness of the intervention suffers as a result. Adaptive control approaches are potential candidates to overcome this problem. However, they have seen limited application to FES-based rehabilitation. Most notably, model reference adaptive control (MRAC) has been applied and tested with able-bodied participants [11]. However, it has not transferred to the clinical domain.

The application of adaptive control approaches to real-world applications relies heavily on robustness, which, in the mid-1980s, was shown to be largely absent [12]. There followed a period of intense interest in establishing a complete robust adaptive control theory [13]. A recent approach has used the framework [14] to enable a robustness analysis

of adaptive control schemes [15]. This framework was a generalisation of the linear gap metric theory that underpins robust control, and enabled the removal in [15] of the restrictions attached to previously used uncertainty models in robust adaptive control [13]. Recent research [16], [17], [18] has employed this framework to provide robustness results for multiple model type adaptive control (MMAC) algorithms [19], [20], [21]. In estimation-based multiple model switched adaptive control (EMMSAC), optimal disturbance estimation is used to assess the performance of a number of pre-selected candidate plant models. A controller design procedure assigns to each plant candidate a corresponding controller such that the controller belonging to the plant with best performance can be switched into closed-loop with the plant. The EMMSAC framework, along with stability and robustness proofs for linear time-invariant (LTI) plants are developed in [16], [17], [18]. It is distinguished from other approaches to MMAC by its integral treatment of robustness and by its axiomatic input-output analysis, which gives significant generality and wide classes of algorithms. Further, the approach gives a principled theoretical route to design. This paper represents the first major implementation of EMMSAC, but note the application to atomic force microscopy in [22]. Also note that the Kalman Filter based optimisation used in this variant of EMMSAC is strongly related to the early (stochastic) versions of MMAC (see e.g. [19]).

This paper comprises the first application of multiple model switched adaptive control to regulate FES-induced isometric muscle contractions, and establishes its potential for stroke rehabilitation. The system comprising the muscle, electrodes, and stimulator is both nonlinear and time-varying; modifications to the EMMSAC framework are hence introduced that allow its implementation, and in-so-doing motivate the extension of the theory presented in [16], [17], [18] to the time-varying setting. A general multiple model adaptive controller is constructed by analysing experimental data taken from a sample of the population (sample A). Based on principles suggested in [16], [17], [18], this data is used to construct a plant model set that is specific to the experimental plant considered. It is then shown that the resulting controller can be applied to multiple participants (sample B) with limited prior identification and no visible degradation in performance. Statistical analysis of the acquired data shows that the ability to switch between controllers results in improved tracking performance when compared with more conventional, non-switched methods. This paper hence comprises a critical step to realising model-based FES control approaches that do not require explicit identification procedures and adapt to uncertain and time-varying conditions.

II. Problem Description

The most widely assumed structure used to represent electrically stimulated muscle is the Hill model [23], which describes the output force as the product of three independent factors associated with the length and velocity of the muscle, and the nonlinear muscle activation dynamics under isometric conditions. The activation dynamics are most commonly represented by a static nonlinearity in series with linear dynamics (a Hammerstein structure). They constitute an important component of the model, since controlled motions are typically smooth and slow so that the effects of inertia, velocity, and series elasticity are small and the isometric behaviour of muscle dominates. The isometric model captures the principle time-varying effects such as fatigue and spasticity.

The popularity of the Hammerstein structure to represent the activation dynamics is due to its correspondence with biophysics: the static nonlinearity represents the Isometric Recruitment Curve (IRC), which is the static gain relation between the stimulus activation level and steady-state output torque or force when the muscle is held at a fixed length; and the linear dynamics represent the muscle contraction dynamics, which combines with the IRC to give the overall torque generated. Another key advantage of this structure is the ease with which model-based controllers can be implemented. The Hammerstein plant model is illustrated in Figure 1.

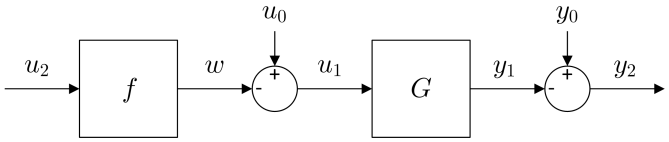

Figure 1. Hammerstein plant model [7], [24]

The Hammerstein structure $P = (f, G)$ is defined by

$$w = f(u_2),\ u_0 = u_1 + w,\ y_1 = Gu_1,\ y_0 = y_1 + y_2$$

in which u_2 is the pulsewidth of the electrical stimulation signal that originates from the stimulator; f is the static nonlinearity, which represents the combined IRC and stimulator dynamics; w is the torque demand (in Nm) after applying f; (u_1, y_1) are the respective input-output signals to the linear block, G, representing the muscle contraction dynamics; (u_0, y_0) are disturbances acting on the system (e.g. measurement noise, voluntary input, physical disturbances); and y_2 is the output torque. The input disturbance, u_0, is assumed to appear between the nonlinear and linear blocks. Note that this ensures that cancellation of the nonlinearity can occur in situations where $u_0 \neq 0$. In general, the disturbances (u_0, y_0), and therefore the plant signals (u_1, y_1), are unknown. The control input, u_2, and the measurable output, y_2, are known signals.

The stimulator output is typically a square wave of fixed amplitude and frequency, as illustrated in Figure 2. Each individual pulse is asymmetrical such that only the positive half-cycle has sufficient voltage amplitude to excite the muscle ($|a_1| >$ excitation threshold $> |a_2|$), yet balanced (zero net) current flow is maintained. The pulse frequency, f_s, is typically between $10\ Hz$ and $100\ Hz$. In this research the stimulation frequency is set to either $40\ Hz$ or $80\ Hz$ and the sample rate, T_s, of the control system is set to $160\ Hz$. The controlled variable, u_2, is the pulsewidth, $0 \leq pw \leq 300\ \mu s$.

The nonlinearity, f, has previously been parameterized using saturation, piecewise linear functions and predefined

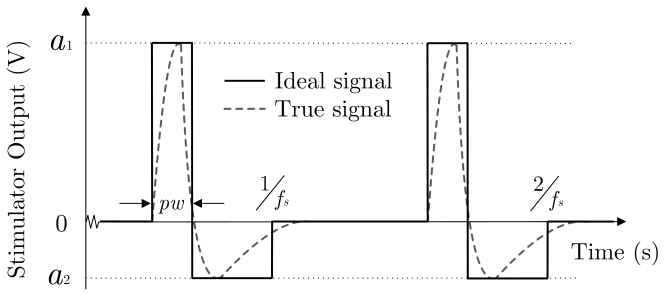

Figure 2. Biphasic stimulation waveform

functional forms (e.g. [25], [26]), with the linear dynamics generally taken as first or second order [27], [28]. In a clinical setting the linear and nonlinear parameters must be identified using an excitation signal that does not cause discomfort for, or an involuntary response from, the participant [24]. Such identification procedures are time-consuming, and accelerate the onset of fatigue and spasticity. The onset of fatigue is further amplified by the inefficiency of the applied stimulation; motor units are all recruited at the same time, in reverse order, and at a higher frequency than normal. Hence, regardless of the initial accuracy of an identified plant model, the fatigue-related time-variance of the true plant will cause a degradation in the performance of the (non-adaptive) control system as time progresses.

III. Estimation-based Multiple Model Switched Adaptive Control

EMMSAC [16], [17], [18] is a multiple model adaptive control approach in which optimal disturbance estimation is used to assess the performance of a set of candidate plant models, $\{P_1, \ldots, P_n\}$. A residual is produced for each plant, which is a measure of the size of its smallest (external) disturbance estimates. The model with smallest residual is deemed to have best performance and its associated controller, chosen from the set of controllers $\{C_1, \ldots, C_n\}$, is switched into closed-loop operation with the true plant.

The closed-loop interconnection of plant P and controller C is given by

$$u_0 = u_1 + u_2,\ y_1 = Pu_1,\ y_0 = y_1 + y_2,\ u_2 = Cy_2 \quad (1)$$

and is illustrated in Figure 3 , where all signals are as defined

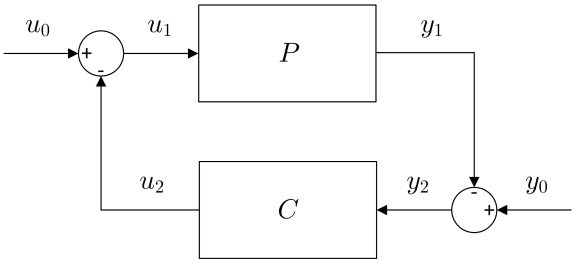

Figure 3. Feedback interconnection of plant and controller

previously in Section II. The gain of the closed-loop $[P, C]$ is given by

$$\gamma = \sup_{(u_0,y_0)\neq 0} \frac{\|u_2, y_2\|}{\|u_0, y_0\|}. \quad (2)$$

If this value is bounded, i.e. if $\gamma < \infty$, then the closed-loop system is said to be BIBO stable. However, in practical implementation, there will be some mismatch between the plant model and the true plant, which motivates the analysis of the chosen control approach from a robust control perspective.

The nonlinear gap metric introduced in [29] allows the robustness analysis of nonlinear systems such as the EMMSAC algorithm. It is shown in [16], [17], [18] that there exist finite bounds from the external disturbances to the internal signals, i.e. $\gamma < \infty$, which guarantees the robustness of the algorithm. An important result is that the established bounds are independent of the size of the plant model set for certain plant set geometries [16], [17], [18].

The sections that follow introduce the general framework that describes EMMSAC for LTI plants and highlight the essential, non-trivial choices that are required to allow its application to the nonlinear, time-varying control problem considered here. To date, EMMSAC has been developed for the time-invariant case only. The work presented hence provides insight into the future progression of theory to the time-varying setting.

A. Disturbance Estimation

For each LTI plant model, $P_p \in \{P_1, \ldots, P_n\}$, the disturbance estimation procedure produces an estimate of the size of the smallest disturbances, (u_0^p, y_0^p), that explain the observed signals, (u_2, y_2). The residual, r_p, which is a measure of the size of the disturbance estimate, is then used to evaluate the performance of the associated plant, where a smaller value represents a model whose dynamics best 'explains' the observations. Note that the estimation can be performed on both an infinite- and finite-horizon [16], [17], [18]; the remainder of this description deals exclusively with the infinite-horizon case.

Let $\mathcal{N}_p^{[a,b]}(u_2, y_2)$ be the set of all disturbance signals (u_0^p, y_0^p) that are compatible with the observed signals (u_2, y_2) and (1) for the candidate plant P_p, over the interval $[a, b],\ a \leq b,\ a, b \in \mathbb{Z}$. Then the residual at time k for plant p is given by

$$r_p[k] = \inf\{r \geq 0 \mid r = \|v_0\|,\ v_0 \in \mathcal{N}_p^{[0,k]}(u_2, y_2)\}.$$

Computation of the optimal disturbance estimates v_0 is, in general, not a recursive procedure. Thus direct computation of $r_p[k]$, in general, will result in a growth in computational complexity for increasing k. However, there is a special case for which the computation of $r_p[k]$ is recursive; when the estimation is performed using a Kalman filter [30] for LTI plants in l^2, as justified by Theorem III.1 below. In that case, the computational complexity of the estimation depends only on the order of the plant model. This result enables a computationally tractable implementation of the infinite-horizon optimal disturbance estimator.

B. Kalman Filter Implementation

The infinite-horizon Kalman filter disturbance estimator is now introduced. Given a candidate linear plant model, P_p, with initial state $x_p(0)$ and discrete-time state space representation given for $k \in \mathbb{N}$,

$$x_p(k+1) = A_p x_p(k) + B_p u_1^p(k),\ y_1^p(k) = C_p x_p(k) \quad (3)$$

where $x_p(0)$ has the possibility of being non-zero, the discrete-time Kalman filter estimator, KF_p, is given by

$$\begin{aligned}
\hat{x}_p(k+1/2) &= \hat{x}_p(k) + \Sigma_p(k) C_p^\top [C_p \Sigma_p(k) C_p^\top + I]^{-1} \cdot \\
&\qquad \cdot [y_2(k) - C_p \hat{x}_p(k)] \\
\Sigma_p(k+1/2) &= \Sigma_p(k) - \Sigma_p(k) C_p^\top [C_p \Sigma_p(k) C_p^\top + I]^{-1} \cdot \\
&\qquad \cdot C_p \Sigma_p(k) \\
\hat{x}_p(k+1) &= A_p \hat{x}_p(k+1/2) + B_p u_2(k) \\
\Sigma_p(k+1) &= A_p \Sigma_p(k+1/2) A_p^\top + B_p B_p^\top \\
\tilde{y}_1^p(k) &= C_p \hat{x}_p(k)
\end{aligned} \quad (4)$$

with model-dependent initial covariance and initial state estimate, $\Sigma_p(0)$ and $\hat{x}_p(0)$, respectively. In these equations, $u_2(k)$ and $y_2(k)$ are as defined previously; and $\tilde{y}_1^p$ is the Kalman filter estimate of the plant output, y_1, before disturbance y_0, associated with model P_p.

In this Kalman filter implementation, the residual is given by the following theorem from [17].

Theorem III.1. *Let $G_p = (A_p, B_p, C_p, 0)$ be the state-space representation of an LTI plant model. Consider the Kalman filter equations (4). Let $\Sigma_p(0) = 0$ and $\hat{x}_p(0) = 0$. Then the residual for plant model G_p at sample instant τ is given by the weighted l^2-norm*

$$r_p[\tau] = \left[\sum_{k=0}^{\tau} \| y_2(k) + \tilde{y}_1^p(k) \|^2_{[C_p \Sigma_p(k) C_p^\top + I]^{-1}} \right]^{1/2}.$$

Proof. The proof can be found in [16]. □

Note that $\Sigma(k) \geq 0$ for all $k \in \mathbb{N}$ if $\Sigma(0) \geq 0$, which guarantees the existence of the weight $[C_p \Sigma_p(k) C_p^\top + I]^{-1}$. The residual calculation can be performed recursively:

$$r_p^2(\tau) = r_p^2(\tau - 1) + \| y_2(\tau) + \tilde{y}_1^p(\tau) \|^2_{[C_p \Sigma_p(\tau) C_p^\top + I]^{-1}}.$$

Note that C_p represents the state-space output matrix and will also later represent the controller for plant P_p. Throughout this paper, the meaning of C_p is clear from the context in which it is used.

Use of a Kalman filter to perform the disturbance estimation limits the algorithm to the LTI setting. To perform the disturbance estimation for the nonlinear Hammerstein plant model, $P_p = (f_p, G_p)$, the nonlinear function, f_p, is first applied to the observed signal u_2 to produce a hypothesis, $\hat{w}_p$, for the internal signal, w, of Figure 1, $(f, G) = (f_p, G_p)$. The estimation is then performed exactly as in the linear case where the estimator (4) is designed for G_p and utilises observed signals $(\hat{w}_p, y_2)$. The same procedure to perform the estimation for linear plant models with input saturation is theoretically justified in [17].

C. Switching Algorithm

Given a set of n plant models, the switching signal points to the plant $P_p,\ 1 \leq p \leq n$ that has minimal residual $r_p[k]$ and index p at sample instant k, and is thus given by

$$q(k) := \underset{1 \leq p \leq n}{\operatorname{argmin}}\ r_p[k],\ \forall k \in \mathbb{N}.$$

If there are multiple minimal residuals $\{r_{p_1}, \ldots, r_{p_b}\}$ at time step k, then $q(k)$ is taken to be the smallest index p_i, $q(k) = \min\{p_1, \ldots, p_b\}$. Once a plant model has been identified, its corresponding controller is switched into the closed-loop with the true plant. The controller is initialised with a zero internal state when it is switched into the closed loop. Thus we define the controller $C : u_2 \mapsto y_2$ by $u_2(k) = C_{q(k)}(v(t))$ where

$$v(t) = \begin{cases} 0 & \text{if } t < k_s \\ y_2(t) & \text{if } t \geq k_s \end{cases}$$

and k_s is the sample at which the last switch occurred, i.e. $k_s = \max\{t \leq k \mid q(t) \neq q(t-1)\}$. This has simple interpretation; at sample instant k_s, controller $C_{q(k)}$ is switched into the closed-loop with zero initial conditions, and remains there until the next sample instant, k, at which the switching signal, $q(k)$, changes value. An alternative approach (not pursued here) is to match the states of the old and new controller to provide a smoother, 'bumpless' transfer at switching times. Additionally, there may exist disturbances, (u_0, y_0), which might cause q to switch rapidly and lead to instability. To prevent this, an operator, D, can be used to assign to each plant model a minimum delay that must elapse before a switch is allowed to occur. Note that this delay operator is not implemented for the practical application presented here.

D. Extension to the Time-Varying Case

The approach presented in the preceding sections is required for the implementation of the algorithm in the time-invariant setting. For the time-varying case considered in this application, the model initially identified to represent the true plant becomes less appropriate as the physical characteristics that define the true plant vary with time. Thus, the information gained from observation of (u_2, y_2) at the start of the horizon has less significance than the information gained from the most recent observations of (u_2, y_2). To account for this, a weight can be applied each time the residual is calculated to successively weight older observations towards zero. Application of the weight, $\lambda < 1$, gives recursive residual calculation

$$r_p^2(\tau) = \lambda r_p^2(\tau - 1) + \| y_2(\tau) + \tilde{y}_1^p(\tau) \|^2_{[C_p \Sigma_p(\tau) C_p^\top + I]^{-1}}.$$

Although it allows improved controller performance when applied to time-varying plants such as those considered in this paper, this modification is yet to be justified theoretically within the EMMSAC framework. We note, however, that the introduction of 'forgetting factors', $\lambda < 1$, is common in adaptive control for time-varying plants.

IV. Application of EMMSAC framework to nonlinear FES muscle models

This section follows on from the problem description given in Section II. The details provided here enable the practical implementation of EMMSAC in Section V.

A. Model Specification

Due to confirmed accuracy in clinical trials [7], in this paper the linear component, G, of the underlying Hammerstein plant model is taken as the critically damped second order transfer function

$$G:\ u_1 \mapsto y_1,\ Y_1(s) = \frac{\omega_n^2}{s^2 + 2\omega_n s + \omega_n^2}\, U_1(s) \qquad (5)$$

in which ω_n, the natural frequency, is the uncertain parameter to be identified. The discrete-time state space representation of this transfer function is obtained by first taking the observable canonical form [31], given by

$$A_o = \begin{bmatrix} -2\omega_n & 1 \\ -\omega_n^2 & 0 \end{bmatrix},\ B_o = \begin{bmatrix} 0 \\ \omega_n^2 \end{bmatrix},\ C_o = \begin{bmatrix} 1 & 0 \end{bmatrix} \qquad (6)$$

and discretising the resulting matrices at the specified sample rate, T_s, to produce for the linear model, G, a corresponding discrete-time state-space representation $(A, B, C, 0)$. The discretisation process uses zero order hold on the inputs. The observable canonical form is used so that $C = \begin{bmatrix} 1 & 0 \end{bmatrix}$ for all ω_n, which simplifies later control design.

The static nonlinearity takes the form

$$f:\ u_2 \mapsto w,\ w = a_1\left(\frac{e^{a_2 u_2} - 1}{e^{a_2 u_2} + a_3}\right),\ 0 \le u_2 \le 300 \qquad (7)$$

with uncertain parameters a_1, a_2 and a_3 [7]. The upper limit on the pulsewidth input, u_2, is an approximate value at which saturation of the nonlinear function, f, occurs. The limits are included to prevent over-stimulation of the triceps and to ensure the existence of a solution to the inverse function used to linearise the response of the plant. See Figure 8 for representative graphs of f. Note that f always has a sigmoidal shape representing a deadband region of low response, and an increasing response level with increasing pulsewidth, until saturation is achieved.

B. Plant Model Set

The uncertain true plant is modelled by the Hammerstein structure illustrated in Figure 1 with uncertain nonlinear isometric recruitment curve, f, and uncertain linear activation dynamics, G, as described in Section IV-A. Assuming that the range of each uncertain model parameter is known (from experimental data, for example), it is possible to specify a continuously parameterised set of models that represents the uncertainty associated with the plant. The sets containing all possible values of the uncertain linear and nonlinear Hammerstein model parameters are given respectively by $U_G = [\omega_n^l, \omega_n^u] \subset \mathbb{R}_+$ and $U_f = \{[a_1^l, a_1^u] \times [a_2^l, a_2^u] \times [a_3^l, a_3^u]\} \subset \mathbb{R}_+^3$, in which l denotes a lower bound and u denotes an upper bound. These sets are identified and sampled (as later described in Section V-C) to produce a linear model set and a nonlinear model set, given respectively by $\Delta_G = \left\{(A_i, B_i, C_i),\ i = 1, \ldots, m \mid (6) \text{ holds},\ \omega_n^i \in U_G\right\}$ and $\Delta_f = \left\{f_j,\ j = 1, \ldots, n \mid (7) \text{ holds},\ (a_1^j, a_2^j, a_3^j) \in U_f\right\}$, in which m and n are the required number of linear and nonlinear models, respectively. The complete plant model set, Δ_s, is then formed from the pairwise combination of the elements of Δ_G and Δ_f (i.e. $\Delta_s = \Delta_G \times \Delta_f$) and contains mn plant models $P_p = (f_j, G_i)$.

C. Controller Design

Recall that for each plant model, $P_p = (f_j, G_i)$, $1 \le j \le n$, $1 \le i \le m$, the controller design procedure is required to produce a controller, $C_p : y_2 \mapsto u_2$, that delivers closed-loop stability, as defined by (2). The underlying controller selected in this research has two components, connected in series: a linearising function (designed for f_j), and a linear quadratic (LQ) optimal tracking controller (designed for G_i).

1) Linearising Controller: Linearisation is achieved through application of the inverse function given by

$$f_j^{-1}:\ u^* \mapsto u_2,\ u_2 = \frac{1}{a_2^j} \ln\left(\frac{a_1^j + a_3^j u^*}{a_1^j - u^*}\right),\ T_a^j \le u^* \le T_b^j$$

where u^* represents the output from the LQ optimal controller and a_1^j, a_2^j and a_3^j are the uncertain parameters associated with the j^{th} nonlinear model. The limits on the inverse function arise from the limits imposed on f in (7). In order to find the values of T_a and T_b for a particular plant, the uncertain parameters for that plant need to be substituted into (7) with $u_2 = 0$ and $u_2 = 300$, respectively. The resulting values of w correspond to the values of T_a and T_b, which are then defined for the nonlinear model in question.

2) LQ Optimal Tracking Control: For the LQ optimal state feedback tracking controller, H, implemented here, a quadratic performance index is used to determine the optimal control input. This index penalises non-zero output errors and control signals and its minimisation leads to a linear state feedback law when applied to linear systems, i.e. for LTI plant model G_i, $H_i : x_p \mapsto u^*$. Note that the state, x_p, is indexed by p because it is the state associated with the p^{th} Hammerstein plant model, $P_p = (f_j, G_i)$. Given the state-space representation (3) where $x_p(0)$, and $A_p = A_i$, $B_p = B_i$ and $C_p = C_i$ are discrete, the performance index is given by

$$J = \sum_{t=t_0+1}^{T} \Big\{[x_p(t) - \tilde{x}(t)]^\top Q [x_p(t) - \tilde{x}(t)] + u^\top(t-1) R u(t-1)\Big\}, \qquad (8)$$

where $\tilde{x}(t)$ is the reference state trajectory, Q is a positive-semidefinite matrix weight used to penalise tracking error, and R is a positive-definite matrix weight used to penalise control effort. Given fixed $\tilde{x}(t)$, Q, and R, the resulting LQ optimal controller, minimises the weighted costs associated with tracking error and control effort. Selection of this controller design procedure allows both the fair comparison of multiple controllers, and the level of FES-based assistance to be set by the therapist when used in a clinical setting (by varying the weights Q and R). For the single-input single-output (SISO) case considered here, R can be taken as unity

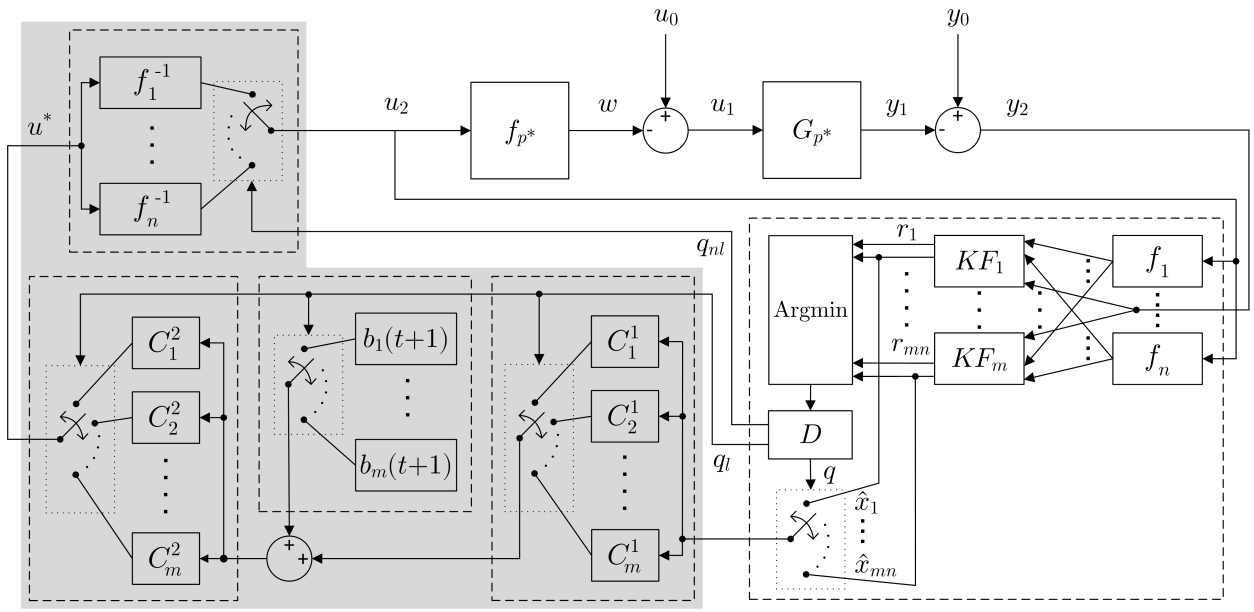

Figure 4. Control system showing uncertain plant, estimator bank, and switched controller

without loss of generality. The matrix Q is a diagonal matrix whose dimension is dependent on the order of the system. In this particular case, the system is assumed to be second order and

$$Q = \begin{bmatrix} Q_1 & 0 \\ 0 & Q_2 \end{bmatrix}$$

where Q_1 and Q_2 penalise the first and second elements of the state vector, respectively.

Given an output reference trajectory, $\tilde{y}_2$, the corresponding state trajectory is given by $\tilde{x}_i = C_i^\top (C_i C_i^\top)^{-1} \tilde{y}_2$. Note that use of the observable canonical form for the state space representation (specifically that $C_i = \begin{bmatrix} 1 & 0 \end{bmatrix}$) gives $\tilde{x}_i = [\tilde{y}_2 \ 0]^\top$. The importance of this is that the reference state trajectory, $\tilde{x}$, is then the same for all linear models, G_i, which simplifies the implementation of the controller.

Given the state-space representation (3), $A_p = A_i$, $B_p = B_i$ and $C_p = C_i$, and the performance index (8), the discrete-time solution for optimal control is given by

$$\begin{aligned} u_i^*[x_p](t) = -[B_i^\top S_i(t+1) B_i + R]^{-1} B_i^\top \cdot \\ [S_i(t+1) A_i x_p(t) + b_i(t+1)] \end{aligned} \tag{9}$$

where $S_i(\cdot)$ and $b_i(\cdot)$ are recursively computed using

$$\begin{aligned} S_i(t) &= A_i^\top \big\{ S_i(t+1) - S_i(t+1) B_i [B_i^\top S_i(t+1) \cdot \\ &\qquad B_i + R]^{-1} B_i^\top S_i(t+1) \big\} A_i + Q \\ b_i(t) &= [A_i^\top + K_i(t) B_i^\top] b_i(t+1) - Q \tilde{x}(t) \\ K_i(t) &= -[B_i^\top S_i(t+1) B_i + R]^{-1} B_i^\top S_i(t+1) A_i \end{aligned} \tag{10}$$

with terminal conditions $S_i(T) = Q(T)$, and $b_i(T) = 0$ [32]. Note that the dependence of the optimal control input on the initial state of the system, $x_p(t_0)$, is implicit in (9); although $x_p(t_0)$ does not appear in (9), it determines the value of $x_p(t)$. Also note that the parameters that determine the optimal control, $S_i(\cdot)$ and $b_i(\cdot)$, are computed backwards in time. An advantage of this setup is that the controller can be switched mid-operation without the need to recompute $S_i(\cdot)$ and $b_i(\cdot)$ with new initial conditions.

The described control approach results in a pulsewidth input $u_2 = f_j^{-1}[u_i^*[\hat{x}_p](t)]$ where $\hat{x}_p$ is the Kalman filter state estimate for plant p obtained from (4) and this defines the p^{th} controller, C_p, corresponding to $P_p = (f_j, G_i)$.

D. Full Control System

The full control system, as described in the preceding sections, is illustrated in Figure 4. The true plant is assumed to have Hammerstein structure with components f_{p^*} and G_{p^*}, and disturbances, (u_0, y_0), as described in Section II. The control input, u_2, is fed into the bank of n nonlinear functions. The outputs from each of the functions are then fed, along with the measured signal y_2, into the bank of m infinite-horizon Kalman filter estimators. This results in mn estimators, E, where estimator E_{ij} performs the estimation for plant model $P_p = (f_j, G_i)$, as described in Section III-B. Each estimator produces for its plant, P_p, $1 \leq p \leq mn$, a state estimate, $\hat{x}_p$, and a residual, r_p, that indicates the size of the smallest estimated disturbances, $(\hat{u}_0^p, \hat{y}_0^p)$, for that plant model. The weighted residual calculation is performed as described in Section III-D. Use of the weight, λ, (to successively weight older observations to zero) reduces the switching delays that occur as the true plant varies with time. However, the estimation becomes more susceptible to noise and rapid, potentially destabilising, switching can occur. Thus, selection of λ is a balance between achieving switching that is both stable and noise independent, and rapid enough to detect the potentially fast time-variance of the true plant.

Comparison of the residuals $\{r_1, \ldots, r_{mn}\}$ allows the switching signal, q, (as described in Section III-C) to be determined. In turn, this selects the controller $C_p = (f_j^{-1}, H_i)$

to be implemented where the indexes i and j are determined from q. Note that for completeness, the delay operator, D, from [16], [17], [18] has been included in the figure; for the implementation described in this paper no switching delay is enforced. The switched controller (grey highlight) is divided into two sections, one containing LQ optimal tracking controllers, the other containing inverse nonlinear functions, as detailed in Section IV-C. Thus within D the switching signal q is decomposed into two additional switching signals: $1 \leq q_l \leq m$, specifying the switched linear quadratic optimal tracking controller; and $1 \leq q_{nl} \leq n$, specifying the switched inverse nonlinear function. The former is used with the switched state estimate $\hat{x}_q$ (observe that q, rather than q_l, switches the state estimate) to implement the switched LQ optimal state feedback controller (9). The LQ optimal tracking controller is separated into three blocks, C_i^1, C_i^2, and $b_i(t+1)$, $1 \leq i \leq m$, where m is the total number of linear models (and hence the total number of LQ optimal tracking controllers), b_i is determined from (10), and where $C_i^1[\hat{x}_p](t) = S_i(t+1)A_i\hat{x}_p(t)$, $C_i^2[v](t) = \{-[B_i'S_i(t+1)B_i + R]^{-1}B_i'\}v(t)$, and $v(t) = C_i^1[\hat{x}_p](t) + b_i(t)$. The resulting control signal, u^*, as defined by (9) is fed into the bank of inverse nonlinear functions, f_j^{-1}, $1 \leq j \leq n$, where n is the total number of nonlinear models. Then the nonlinear switching signal, q_{nl}, specifies the switched inverse nonlinear function that maps from u^*, to the pulsewidth, u_2, of the electrical stimulus.

V. Practical Application

This section builds on the basis of Section III, and the application-specific design framework of Section IV, to develop a principled experimental procedure for constructing the plant model set for a sample (A) of the population. This set is then used in tests on another sample (B) of the population, with experimental results given in Section VI.

A. Experiment Apparatus

The experimental facility comprises a workstation that has been used in both clinical stroke rehabilitation trials [8] and model identification tests [24], [33]. The system is validated in [34] and employs a planar robotic arm that supports users in performing two-dimensional reaching tasks. A six-axis force-moment sensor is attached to the end effector and in this research the positions of both the shoulder and end effector are fixed such that the angle between the ulna and humerus is approximately ninety degrees, as shown in Figure 5.

Surface electrodes (PALS Platinum 5×5cm) are placed on the triceps such that the induced currents flow through muscle fibres in the lateral and long heads to achieve the strongest possible contraction. The surface electrodes are attached to the output channel of a commercial stimulator (Odstock, O4CHS). The pulsewidth is the controlled variable with the pulse frequency initially fixed at 80 Hz (see Figure 2 for an illustration of the waveform). For each participant, the amplitude is set at a comfortable level by applying a maximum pulsewidth input ($300\ \mu s$) and then gradually increasing the stimulator amplitude.

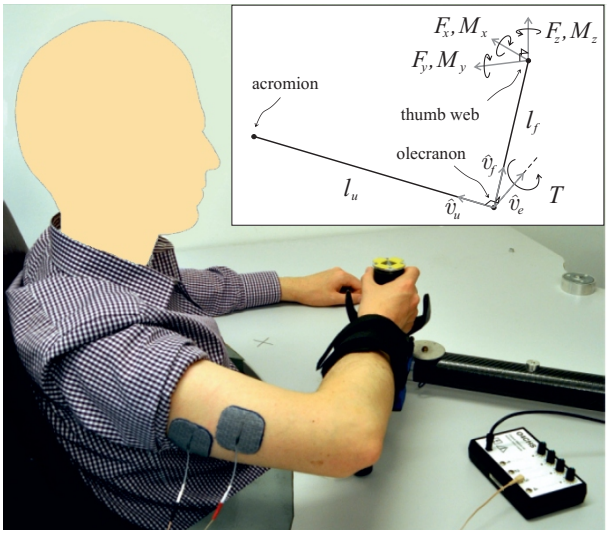

Figure 5. Participant performing isometric task using robotic workstation

The clinically relevant task is for the torque, T, produced about the elbow to follow a trajectory that is representative of a smooth reaching movement. The torque is calculated using the kinematic model shown in Figure 5, where F_x, F_y, and F_z are the force components measured by the sensor in the base coordinate frame; M_x, M_y, and M_z are the moment components, also measured by the sensor, also in the base coordinate frame; l_f is the distance from the thumb web to the olecranon process of the ulna; l_u is the distance from the olecranon process to the acromion; $\hat{v}_f$ and $\hat{v}_u$ are unit vectors aligned with the forearm and upper arm, respectively; and $\hat{v}_e = \hat{v}_f \times \hat{v}_u$ is the axis about which the stimulated triceps exerts a moment. The torque amplitude is calculated from $T = \begin{bmatrix} F_x & F_y & F_z \end{bmatrix}^\top \cdot (\hat{v}_e \times \hat{v}_f) \cdot l_1 + \begin{bmatrix} m_x & m_y & m_z \end{bmatrix}^\top \cdot \hat{v}_e$. To identify a kinematic model, the lengths l_f and l_u are directly measured, and the approximate acromion position is calculated via a non-linear optimisation procedure applied to range of movement data in which the participant moves their arm over the workspace [35]. The vectors, $\hat{v}_f$ and $\hat{v}_u$, are then calculated from the known end-effector position and the underlying kinematic relationship.

B. Stimulated Muscle Model Identification Procedure

To produce a realistic set of plant candidates, an identification procedure was used to produce an estimate of the uncertain Hammerstein parameters from experimental data. The ramp deconvolution method is commonly used to achieve this and has been shown to have high accuracy with relatively short computation time when compared to other identification methods [28]. Results are presented in [24], which show that the method is comparable to other tests used in the literature.

To identify a single model representing the true plant, the participant was seated at the workstation with their forearm placed in the support as in Figure 5, and the pre-experiment set up was performed as described in Section V-A. Next a trapezoidal input signal was provided to the stimulator and the resulting pulsed stimulation waveform applied to the muscle.

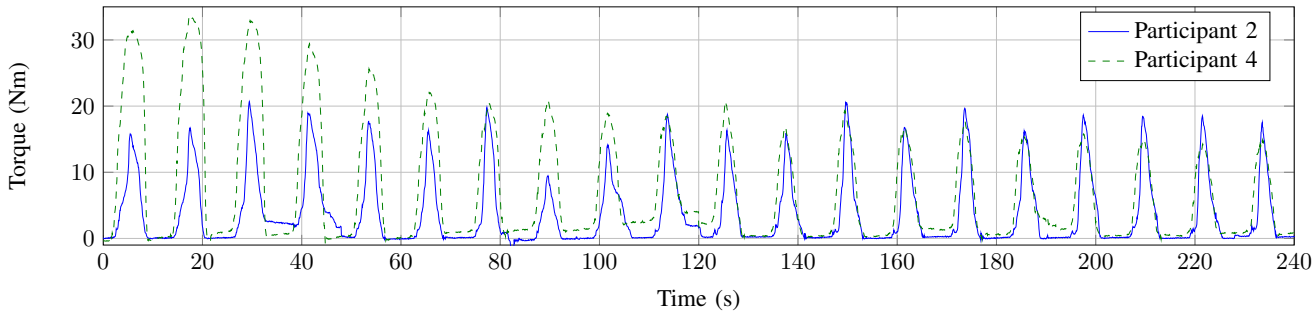

Figure 6. Comparison of torque output for two participants

The input (u_2) is 0 μs for five seconds, ramps up to 300 μs over five seconds, ramps back down to 0 μs over five seconds, and remains at 0 μs for a further five seconds. The zero stimulation period allows the triceps to relax gradually, to ensure that the muscle returns to the same initial (resting) conditions for each application of the reference. After filtering the recorded response (third order lowpass Butterworth, cut-off frequency 10 Hz), the curve is shifted to minimise the difference between it and the reference input on the assumption that the linear activation dynamics comprise a pure delay, which allows the nonlinear and linear dynamics to be deconvolved. This leaves an approximation of signal $w = u_1$ (assuming zero disturbance, i.e. $u_0 = 0$). This can be plotted against u_2 (the reference input) to directly obtain the shape of f (since u_2 comprises a straight line from 0 to 300 and back to 0). Using this information, a least squares nonlinear curve fitting procedure (MATLAB function *lsqcurvefit*) is employed, which identifies the optimal (in a least squares sense) nonlinear parameters of (7). Utilising the identified nonlinearity to form the input w, an iterative procedure is then used to identify the linear parameter of (5) and (6): for each possible value of ω_n, the simulated response (using ω_n and the identified nonlinear parameters, (a_1, a_2, a_3)) is compared with the true response. The optimal value of ω_n is deemed to be that which minimises the error between these two signals. The fitting process is validated by comparing the modelled response with the true response.

C. Selection of Plant Model Set

The selection of an appropriate plant model set is required to ensure a high level of performance for all possible true plants. This section details a suitable procedure to identify plant candidates for the population (young, able bodied participants) using experimental data.

1) Collection of Fatigue Data: Five able-bodied, young participants were recruited for the initial stage of data collection (sample A). Participants were seated at the workstation with their arm placed in the forearm support as shown in Figure 5, and the setup procedure described in Section V-A was performed. For two of the participants it was found that optimising the electrode position resulted in the early onset of fatigue for a single arm. As a result, the test data (both participants, single arm only) was discounted.

To identify plant models that capture the response of the triceps to electrical stimulation at different stages of fatigue, the ramp deconvolution input signal described in Section V-B was applied repeatedly to the triceps of each participant. Each application of this reference signal is referred to as a *trial*, and thirty consecutive trials are referred to as a *fatigue test*. The developed torque recorded from two different participants engaging in such a test is shown in Figure 6, which illustrates the inter-participant and inter-trial variation. Upon completion of the fatigue test, the identification procedure described in Section V-B was used to identify thirty plant models (one for each trial) that captured the time-variance of the true plant. The complete data set contained 240 plant models (eight fatigue tests in total, thirty models identified in each).

2) Analysis of Fatigue Data: Table I shows for each participant the correlation between the identified value of the linear parameter, ω_n, and the trial number. The 'p-value' indicates the probability of obtaining the result when the true correlation is in fact zero. Although there is some suggestion of an overall decrease in ω_n with increasing trial number (average correlation $= -0.31$), this is not consistently observed when the data sets are considered individually. Thus it cannot be concluded that any correlation between the two parameters exists. Lack of correlation implies that the trial number cannot be used as a scheduling variable to predict the linear parameter in advance. Figure 7 shows the trial-by-trial variation of the identified linear parameter, ω_n, for the eight fatigue tests performed. Observe that for all

Participant	**Arm**	$-1 \leq$ **Correlation** ≤ 1	**p-value**
1	L	-0.72	0.0000
	R	-	-
2	L	0.52	0.0035
	R	-0.34	0.0643
3	L	0.09	0.6539
	R	0.10	0.5939
4	L	-0.40	0.0299
	R	-	-
5	L	-0.58	0.0009
	R	0.25	0.1778
Average	-	-0.31	0.0998

Table I
CORRELATION BETWEEN LINEAR PARAMETER AND TRIAL NUMBER

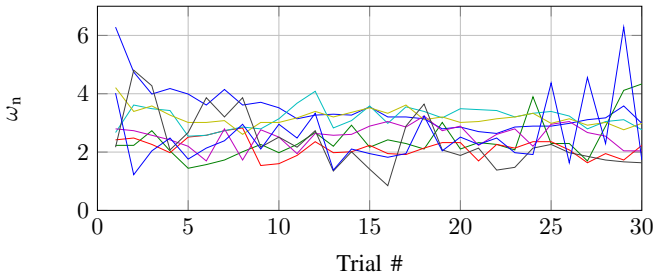

Figure 7. Variation of linear parameter, ω_n for all tests on all participants

participants the majority of the identified values lie within a narrow band. Thus a single, relatively small parameter set can be chosen to represent the uncertain linear activation dynamics for the population. This is a significant outcome, as smaller model sets result in reduced computational costs. It is possible to parameterise the linear parameter set using the cumulative distribution, Φ_{ω_n}, of the data shown in Figure 7. This distribution represents the linear parameter uncertainty set, U_G, of Section IV-B, which can be sampled to produce the linear model set, Δ_G. The i^{th} value of ω_n in Δ_G is given by $\omega_n^i = \Phi_{\omega_n}^{-1}((2i-1)/2m)$, $1 \leq i \leq m$, where, as previously, m is the required number of linear models.

The same procedure cannot be applied to identify the uncertainty set representing the nonlinear component of the model because it is described by three parameters, (a_1, a_2, a_3), which only have partial correlation with one another and the trial number. However, a few general observations can be made regarding the variation of the identified nonlinearity. As the trial number increases: i) it takes a larger initial pulsewidth to achieve a contraction, resulting in a larger *deadband* region; ii) for high values of pulsewidth the output saturates less, resulting in a smaller *saturation* region; and iii) the peak torque produced (at a pulsewidth input of 300 μs) decreases. These observations are illustrated in Figure 8, which shows the nonlinear models identified for a single fatigue test for one participant (tr_j indicates the isometric recruitment curve identified using data from the j^{th} trial). Given these observations, selection of the nonlinearity follows a similar procedure to selection of the linear model. First, the cumulative distributions (Φ_{db} and Φ_{sat}, respectively) of the upper deadband and lower saturation pulsewidth limits for all 240 identified models are found. Here the deadband upper pulsewidth limit and the saturation lower pulsewidth limit are defined as the pulsewidths at which the output of the nonlinearity reaches 5% and 95% of its peak value, respectively. The distribution of the peak torque is not found using experimental data as the required peak torque is task- and participant-dependent and varies as the stimulator gain is tuned. For the tests described here, the largest peak torque output was selected to be 20 Nm (corresponding to the strongest non-fatigued participant); and the smallest peak torque output was set to 2 Nm (corresponding to the weakest fatigued participant). The cumulative distribution, $\Phi_{T_{peak}}$, of the peak torque output value was then found by interpolating linearly between these upper and lower peak torque limits. The three cumulative distributions, Φ_{db}, Φ_{sat}, $\Phi_{T_{peak}}$, are then used to find the set of parameters (deadband pulsewidth, pw_{db}; saturation pulsewidth, pw_{sat}; and peak output, T_{peak}) for the j^{th} model in Δ_f using $pw_{db}^j = \Phi_{db}^{-1}((2j-1)/2n)$, $pw_{sat}^j = \Phi_{sat}^{-1}((2j-1)/2n)$, and $T_{peak}^j = \Phi_{T_{peak}}^{-1}((2j-1)/2n)$, $1 \leq j \leq n$, where n is the required number of nonlinear models. Then an iterative procedure is used to find the (a_1^j, a_2^j, a_3^j) that produce an isometric recruitment curve with deadband, saturation, and peak output that are closest to the required values.

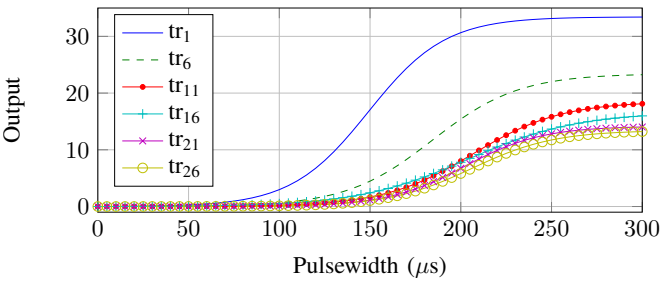

Figure 8. Example trial-by-trial variation of identified nonlinearity

Once the m linear models and n nonlinear models have been identified, all of the pairwise combinations of these sets Δ_G and Δ_f are used to populate the plant model set, which contains mn models.

VI. Experimental Results

A. Experiment Description

Participants were seated at the workstation and the experiment set up of Section V-A was performed. For the controller tests an offset sinusoidal signal of period 10 s was selected as the reference to represent a smooth reaching movement. The (participant-dependent) peak value of this reference was set to 50% of the peak torque measured when maximum stimulation (300 μs) was applied to the triceps. The sample frequency was set to 160 Hz. The stimulator frequency was reduced to 40 Hz to reduce the risk of stimulating peripheral nerves in participants with smaller triceps.

The tracking performance of two different controllers (one adaptive, one non adaptive) was compared. Five able-bodied, young participants were recruited (sample B). Both their left and right arms were tested. Nine data sets are considered hereafter; one data set was discounted because an electrode position that produced an appropriate response could not be found. The only identification tests performed were the measurements of the arm lengths and joint positions used to form a kinematic model of the arm, and the single ramp deconvolution test used to identify a model for use with the non-switched (non-adaptive) controller. For each test, the selected controller was required to regulate the pulsewidth input such that the plant followed the reference torque trajectory over ten consecutive trials. All estimator variables (state and covariance estimates), excluding the residuals, were reset in between trials. After the ten trials were completed, the participant was given a thirty minute rest period to allow

the arm to recover from fatigue before the other controller was tested. The order in which the controllers were applied was randomised to prevent confounding of results.

1) Non-switched controller: The controller used for benchmarking was a single LQ optimal controller with linearising function f^{-1}, designed as described in Section IV-C with controller weights $Q_1 = 50$, $Q_2 = 0$ and $R = 0.4$. The plant model used to design this controller was identified immediately prior to its use using the ramp deconvolution identification procedure described in Section V-B.

2) Switched (adaptive) controller: The switched controller was the estimation-based multiple model switched adaptive controller (EMMSAC) described in Section III. The plant model set was selected a priori using the procedure described in Section V-C with $m = 2$ and $n = 1000$. The estimator for each plant model was the Kalman Filter estimator described in Section III-B; the residual calculation was weighted with a forgetting factor, $\lambda = 0.995$. A linearising controller, as described in Section IV-C, was designed for each plant model $P_p = (f_j, G_i)$ with weights $Q_1 = 50$, $Q_2 = 0$ and $R = 0.4$.

B. Experiment Results

Prior to performing the controller tests the adaptive controller was first tested on a single participant from the population sample used to select the plant model set (sample A of Section V-C). The results of this test are shown in Figure 9. They illustrate the high level of tracking performance that is possible given appropriate tuning of the stimulator gain and an appropriate number of models in the plant model set.

Figure 10 shows the response for a single participant from sample B when performing the test using the non-switched controller. Figure 11 shows results from the same test when using the switched (adaptive) controller. As can be seen from Figure 12, the tracking performance for each controller is comparable over the first five trials. The performance of

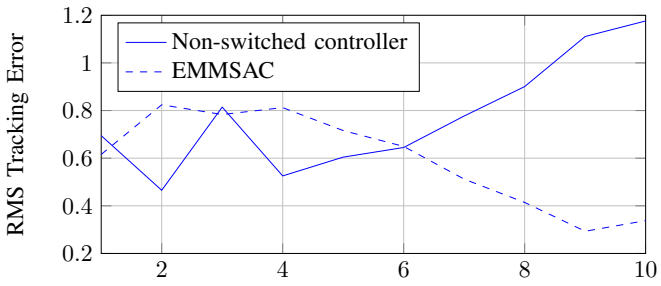

Figure 12. Controller tracking performance

the non-switched controller then degrades as the participant fatigues and the true plant deviates from the model identified prior to the first trial. In contrast, the performance of the switched controller improves over the last five trials. This occurs because the participant relaxes back into a position where approximately zero torque is measured. Note that the lack of complete relaxation on the part of the participant (as seen in the first six trials of Figure 11, in which the response does not accurately track the 'troughs' of the reference) cannot be compensated for by the controller without activating antagonist muscle groups.

The onset of fatigue can be observed from the nonlinear switch shown in Figure 11. The nonlinear models are ordered to represent increasing levels of fatigue from model 1 (the least fatigued) to model 1000 (the most fatigued). Thus, as the participant fatigues, the value of the nonlinear switching signal is expected to increase, which is the trend over all ten trials. This is consistent with the trend observed in Figure 8, which shows the effect of fatigue on the identified nonlinear model. The intra-trial variation of the nonlinear switching signal serves to compensate for minor differences between the reference and response that result from plant-model mismatch, external disturbances, and hysteresis. The latter occurs due to the differences between the FES-activated muscle dynamics and the relaxing muscle dynamics.

During some of the trials (e.g. trial one of Figure 11, after approximately six seconds) the response becomes oscillatory. In this particular example the oscillations are low amplitude. However, in other cases the amplitude has been high enough to significantly reduce the tracking performance. This oscillation occurs because the resistance between the electrodes and the muscle fibres is high, resulting in an 'on-off' mode of operation where the pulsewidth alternates between its minimum and maximum limits (0 μs and 300 μs, respectively, as given by (7)), as the response oscillates about the reference. This can be overcome by replacing the electrodes, the performance (measured by charge transference) of which reduces with increased use. More commonly, the stimulator gain is increased such that an appropriately high current can be induced in the tissues for a larger range of pulsewidths below the maximum value (300 μs). This has been tested experimentally with a subsequent reduction in oscillations. The effect of varying the stimulator gain is to vary the value of the nonlinear switch (increasing the gain decreases the value). This highlights the importance of tuning the stimulator gain when initialising the experiment, such that the true plant lies within the boundaries of the chosen plant model set.

Figure 13 compares the RMS tracking error achieved for the two controllers for all nine sets of data. The introduction

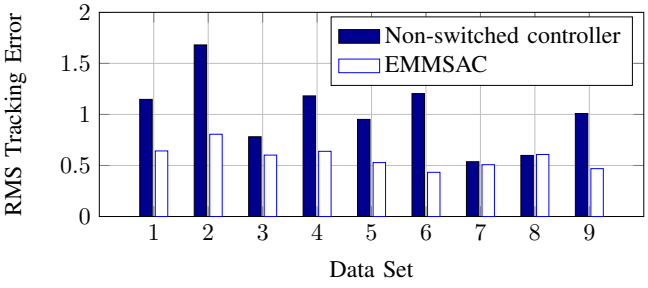

Figure 13. Comparison of controller performance

of switching results in a comparable or reduced level of RMS tracking error for all data sets.

The described experiment is a repeated measures test, which means the two sample sets can be classed as dependent

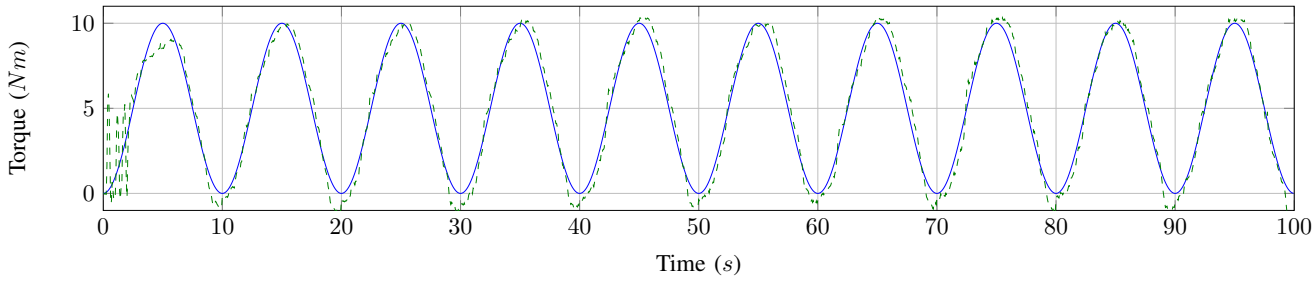

Figure 9. Tracking example, adaptive controller, participant from sample A (solid line represents the reference, dashed line represents the response)

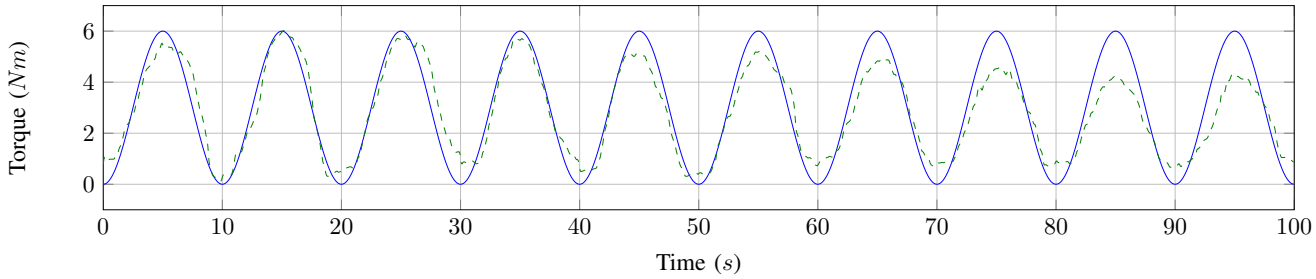

Figure 10. LQ Optimal (non-switched) tracking example, participant from sample B (legend as above)

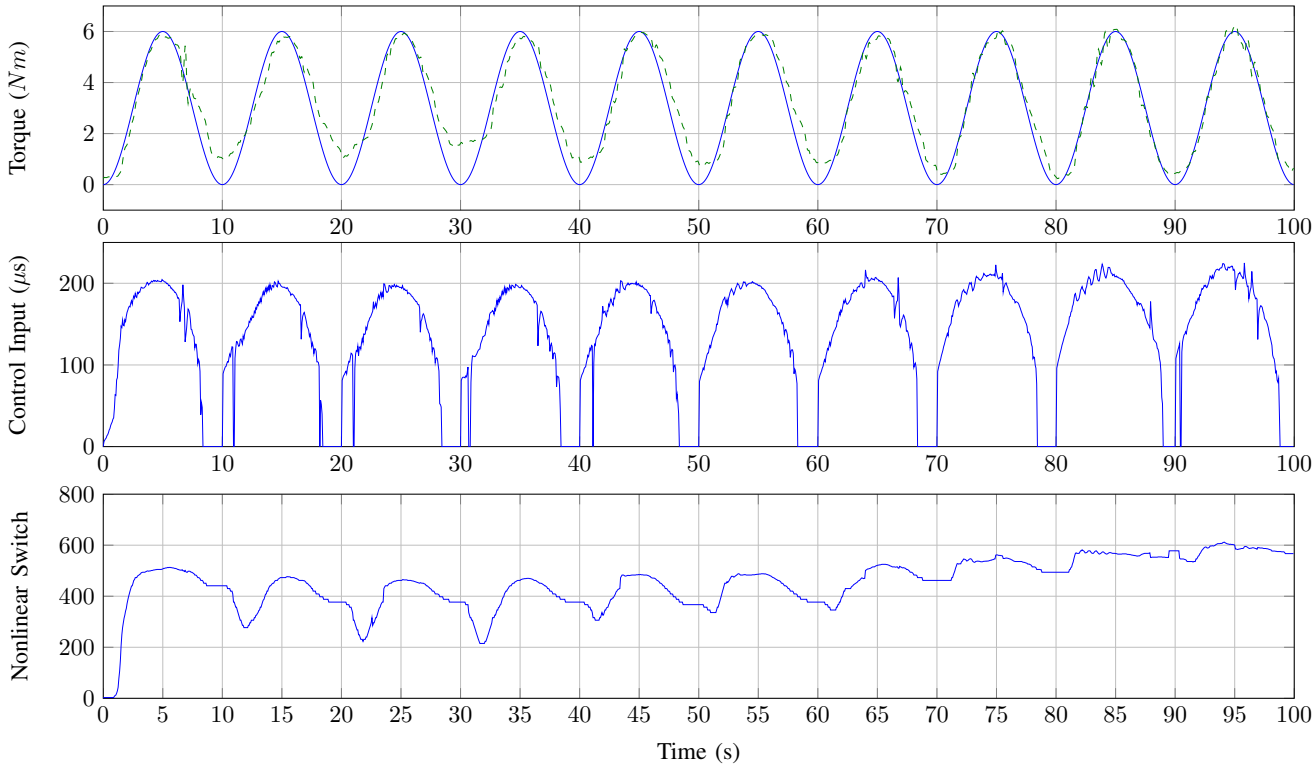

Figure 11. Adaptive controller example, participant from sample B (legend as above)

samples. Thus statistical analysis involves inference about two means with dependent samples. It is assumed that the RMSE data acquired for each of the controllers belongs to a normal distribution. The claim is that the inclusion of switching reduces RMS tracking error. Thus a one-tailed Student's t-test is implemented using the MATLAB function *tcdf()*. For the data shown in Figure 13 there is $> 99\%$ confidence that the inclusion of switching reduces RMS tracking error. At the 95% confidence level the effect of the inclusion of switching is to reduce the average RMS tracking error by approximately 22%. If the second half of the data (trials five to ten, at which point fatigue is more prevalent) are

considered alone, the tracking error is reduced by 42% at the 95% confidence level. This is a significant result. Note that performance gains resulting from the inclusion of switching are likely to be higher still in a clinical setting because stroke patients are likely to experience higher levels of fatigue (and spasticity) more rapidly than able-bodied participants.

VII. Conclusions and Future Work

The results presented in this paper confirm the ability of the multiple model switched adaptive controller to compensate for the time-varying characteristics of human muscle. Furthermore, the hypothesis that a single plant model set can be used for the entire population of able-bodied, young participants has been confirmed; combined with reduced tracking error, the ability to apply a controller with minimal prior identification tests is an attractive characteristic of the proposed algorithm, as it potentially enables the transference of closed-loop model-based control approaches to the clinical setting. This, in turn, will allow more accurate assistance and effectiveness of rehabilitation, which will enable the efficacy of FES as a treatment to be tested. Future work will involve: testing the ability of the adaptive controller to track randomised smooth reference signals; extension to the non-isometric case in which the position of the electrodes changes relative to the muscle; inclusion and modelling of voluntary effort from participants; and testing with the populations of older, able-bodied participants, and stroke patients.

References

[1] World Health Organisation, "The Atlas of Heart Disease and Stroke, Part three: the burden, Section 15: Global burden of stroke," p. 50.

[2] ISWP, "National clinical guideline for stroke," Royal College of Physicians, Tech. Rep., 2012.

[3] J. H. Burridge and M. Ladouceur, "Clinical and Therapeutic Applications of Neuromuscular Stimulation: A Review of Current Use and Speculation into Future Developments," *Neuromodulation*, vol. 4, no. 4, pp. 147–154, 2001.

[4] J. R. de Kroon, J. H. van der Lee, M. J. IJzerman, and G. J. Lankhorst, "Therapeutic electrical stimulation to improve motor control and functional abilities of the upper extremity after stroke: a systematic review," *Clinical Rehabilitation*, vol. 16, no. 4, pp. 350–360, 2002.

[5] J. R. de Kroon, M. J. Ijzerman, J. Chae, G. J. Lankhorst, and G. Zilvold, "Relation between stimulation characteristics and clinical outcome in studies using electrical stimulation to improve motor control of the upper extremity in stroke," *Journal of rehabilitation medicine*, vol. 37, no. 2, pp. 65–74, 2005.

[6] D. N. Rushton, "Functional Electrical Stimulation and rehabilitation - an hypothesis," *Medical Engineering & Physics*, vol. 25, no. 1, pp. 75–78, 2003.

[7] C. T. Freeman, A.-M. Hughes, J. H. Burridge, P. H. Chappell, P. L. Lewin, and E. Rogers, "Iterative learning control of FES applied to the upper extremity for rehabilitation," *Control Engineering Practice*, vol. 17, no. 3, pp. 368–381, 2009.

[8] A.-M. Hughes, C. T. Freeman, J. H. Burridge, P. H. Chappell, P. L. Lewin, and E. Rogers, "Feasibility of Iterative Learning Control Mediated by Functional Electrical Stimulation for Reaching After Stroke," *Neurorehabilitation and Neural Repair*, vol. 23, no. 6, pp. 559–568, 2009.

[9] K. L. Meadmore, A.-M. Hughes, C. T. Freeman, Z. Cai, D. Tong, J. H. Burridge, and E. Rogers, "Functional electrical stimulation mediated by iterative learning control and 3D robotics reduces motor impairment in chronic stroke," *Journal of NeuroEngineering and Rehabilitation*, vol. 9, no. 32, 2012.

[10] C. L. Lynch and M. R. Popović, "Functional Electrical Stimulation: Closed-Loop Control of Induced Muscle Contractions," *IEEE Control Systems Magazine*, pp. 40–50, 2008.

[11] J. Allin and G. F. Inbar, "FNS Parameter Selection and Upper Limb Characterization," *IEEE Transactions on Biomedical Engineering*, vol. 33, no. 9, pp. 809–817, 1986.

[12] C. Rohrs, L. Valavani, M. Athans, and G. Stein, "Robustness of Continuous-Time Adaptive Control Algorithms in the Presence of Unmodeled Dynamics," *IEEE Transactions on Automatic Control*, vol. 30, no. 9, pp. 881–889, 1985.

[13] P. Ioannou and J. Sun, *Robust Adaptive Control*. Prentice Hall, 1996.

[14] T. T. Georgiou and M. C. Smith, "Robustness Analysis of Nonlinear Feedback Systems: An Input-Output Approach," *IEEE Transactions on Automatic Control*, vol. 42, no. 9, pp. 1200–1221, 1997.

[15] M. French, "Adaptive Control and Robustness in the Gap Metric," *IEEE Transactions on Automatic Control*, vol. 53, no. 2, pp. 461–478, 2008.

[16] D. Buchstaller, "Robust Stability and Performance for Multiple Model Switched Adaptive Control," Ph.D. dissertation, University of Southampton, 2010.

[17] D. Buchstaller and M. French, "Robust stability and performance for multiple model adaptive control: Part i - the framework (provisionally accepted)," *IEEE Transactions on Automatic Control*, 2014.

[18] ——, "Robust stability and performance for multiple model adaptive control: Part ii - gain bounds (provisionally accepted)," *IEEE Transactions on Automatic Control*, 2014.

[19] S. Fekri, M. Athans, and A. Pascoal, "Issues, progress and new results in robust adaptive control," *International Journal of Adaptive Control and Signal Processing*, vol. 20, pp. 519–579, 2006.

[20] A. S. Morse, "Supervisory Control of Families of Linear Set-Point Controllers - Part 1: Exact Matching," *IEEE Transactions on Automatic Control*, vol. 41, no. 10, pp. 1413–1431, 1996.

[21] ——, "Supervisory Control of Families of Linear Set-Point Controllers - Part 2: Robustness," *IEEE Transactions on Automatic Control*, vol. 42, no. 11, pp. 1500–1515, 1997.

[22] U. Khan, H. Chong, and M. French, "Automated Controller Tuning for Atomic Force Microscopes Using Estimation Based Multiple Model Switched Adaptive Control," in *52nd IEEE Conference on Decision and Control, Florence, IT*, 2013.

[23] A. V. Hill, "The Heat of Shortening and the Dynamic Constants of Muscle," *Proceedings of the Royal Society of London. Series B, Biological Sciences*, vol. 126, pp. 136–195, 1938.

[24] F. Le, I. Markovsky, C. T. Freeman, and E. Rogers, "Identification of electrically stimulated muscle models of stroke patients," *Control Engineering Practice*, vol. 18, no. 4, pp. 396–407, 2010.

[25] K. J. Hunt, M. Munih, N. N. Donaldson, and F. M. Barr, "Investigation of the Hammerstein Hypothesis in the Modeling of Electrically Stimulated Muscle," *IEEE Transactions on Biomedical Engineering*, vol. 45, no. 8, pp. 998–1009, 1998.

[26] R. Riener and J. Quintern, "A physiologically based model of muscle activation verified by electrical stimulation," *Bioelectrochemistry and Bioenergetics*, vol. 43, pp. 257–264, 1997.

[27] N. Lan, "Stability Analysis for Postural Control in a Two-Joint Limb System," *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, vol. 10, no. 4, pp. 249–259, 2002.

[28] W. K. Durfee and K. E. MacLean, "Methods for Estimating Isometric Recruitment Curves of Electrically Stimulated Muscle," *IEEE Transactions on Biomedical Engineering*, vol. 36, no. 7, pp. 654–667, 1989.

[29] T. T. Georgiou and M. C. Smith, "Optimal Robustness in the Gap Metric," *IEEE Transactions on Automatic Control*, vol. 35, no. 6, pp. 673–686, 1990.

[30] R. E. Kalman, "A New Approach to Linear Filtering and Prediction Problems," *Transactions of the ASME-Journal of Basic Engineering*, vol. 82, no. Series D, pp. 35–45, 1960.

[31] K. Ogata, *Discrete-Time Control Systems*. Pearson Education, 1995.

[32] B. D. O. Anderson and J. B. Moore, *Optimal Control: Linear Quadratic Methods*. Dover Publications, 2007.

[33] F. Le, I. Markovsky, C. T. Freeman, and E. Rogers, "Recursive identification of Hammerstein systems with application to electrically stimulated muscle," *Control Engineering Practice*, vol. 20, no. 4, pp. 386–396, 2012.

[34] C. T. Freeman, A.-M. Hughes, J. H. Burridge, P. H. Chappell, P. L. Lewin, and E. Rogers, "A robotic workstation for stroke rehabilitation of the upper extremity using FES," *Medical Engineering & Physics*, vol. 31, no. 3, pp. 364–373, 2008.

[35] ——, "A Model of the Upper Extremity Using FES for Stroke Rehabilitation," *Journal of Biomechanical Engineering*, vol. 131, no. 3, 2009.