

Motor Unit Number Estimation via Sequential Monte Carlo

Simon A. C. Taylor^{a,*}, Chris Sherlock^a, Gareth Ridall^a, Paul Fearnhead^a

^a*Department of Mathematics and Statistics, Lancaster University, Lancaster LA1 4YF, U.K.*

Abstract

A change in the number of motor units that operate a particular muscle is an important indicator for the progress of a neuromuscular disease and the efficacy of a therapy. Inference for realistic statistical models of the typical data produced when testing muscle function is difficult, and estimating the number of motor units is an ongoing statistical challenge. We consider a set of models for the data, each with a different number of working motor units, and present a novel method for Bayesian inference based on sequential Monte Carlo. This provides estimates of the marginal likelihood and, hence, a posterior probability for each model. Implementing this approach in practice requires a sequential Monte Carlo method that has excellent computational and Monte Carlo properties. We achieve this by benefiting from the model's conditional independence structure, where, given knowledge of which motor units fired as a result of a particular stimulus, parameters that specify the size of each unit's response are independent of the parameters defining the probability that a unit will respond at all. The scalability of our methodology relies on the natural conjugacy structure that we create for the former and an enforced, approximate, conjugate structure for the latter. A simulation study demonstrates the accuracy of our method, and inferences are consistent across two different datasets arising from the same rat tibial muscle.

Keywords: Motor unit number estimation, Sequential Monte Carlo, Model selection, Particle learning

1. Introduction

Motor unit number estimation (MUNE) is a continuing challenge for clinical neurologists. Determination of the number of motor units (MUs) that operate a particular muscle provides important insights into the progression of various neuromuscular ailments such as amyotrophic lateral sclerosis (Shefner et al.,

*Corresponding author.

Email address: s.taylor2@lancaster.ac.uk (Simon A. C. Taylor)

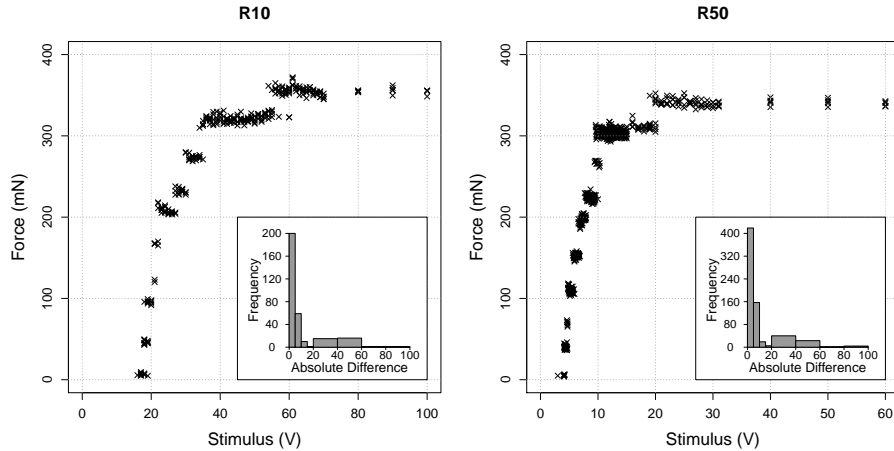


Figure 1: Stimulus-response curve from a rat tibial muscle using $10 \mu\text{sec}$ (left) and $50 \mu\text{sec}$ duration stimuli. Histogram inserts represent the frequency in the absolute difference of twitch forces when ordered by stimulus.

2006; Bromberg, 2007), and aids the assessment of the efficacy of potential therapy treatments (Casella et al., 2010).

A MU is the fundamental component of the neuromuscular system and consists of a single motor neuron and the muscle fibres whose contraction it governs. Restriction to a MU's operation may be a result of impaired communication between the motor neuron and muscle fibres, or abnormality in their function. A direct investigation into the number of MUs via a biopsy, for example, is not helpful since this only determines the presence of each MU, not its functionality.

An electromyography (EMG) investigation into a muscle's function consists of applying a set of electrical stimuli of varying intensity to a group of motor neurons. Each stimulus artificially induces a twitch in the targeted muscle to provide an *in situ* measurement of functioning MUs. The effect on the muscle may be measured by recording either the minute variation in muscle membrane potential or the physical force the muscle exerts (Major and Jones, 2005). The methods developed in this article are applicable to either type of measurement. Since our data consist of whole muscle twitch force (WMTF) measurements we henceforth describe the response in these terms. In a healthy subject, the stimulus-response curve is typically sigmoidal (Henderson et al., 2006), illustrating the smooth recruitment of additional MUs as the stimulus increases; however, the relatively low number of MUs in a patient with impaired muscle function may manifest within the stimulus-response relationship through large jumps in WMTF measurements.

Figure 1 shows the two datasets that will be described and analysed in detail in Section 5, with the large jumps clearly visible. The histograms of absolute differences in response for ordered stimuli show two main modes: one, near 0 mN, corresponding to noise and the other, around 40 mN indicating that

different MUs fired. The noise arises primarily because of small variations in the contribution to the WMTF provided by any MU when it fires. The noise that is visible in isolation at very low stimuli when no MUs are firing is called the baseline noise. This arises from respiration movements and pulse pressure waves, and particular care is taken to minimise such influences, for example by earthing the subject and equipment, restraining the limb, digitally resetting the force signal prior to each stimulus, synchronising stimuli with the pulse cycle and using highly sensitive measurement devices.

MUNE uses the observed stimulus-response pattern to estimate the number of functioning MUs. Techniques for MUNE generally form two classes: the average and comprehensive approaches. The most common averaging approach is the incremental technique of McComas et al. (1971), which assumes that the MUs can be characterised by an ‘average’ MU with a particular single motor unit twitch force (MUTF), estimated as the average of the magnitudes of the observed stepped increases in twitch force. A large stimulus, known as the supramaximal stimulus, is applied in order to cause all MUs to react. The quotient of the WMTF arising from the supramaximal stimulus and the average MUTF provides a count estimate. However, there is no guarantee that a particular single-stepped increase in response corresponds to a new, previously latent, MU, since it may instead be due to a phenomenon called alternation (Brown and Milner-Brown, 1976). This occurs when two or more MUs have similar activation thresholds such that different combinations of MUs may fire in reaction to two identical stimuli. Consequently, the incremental technique tends to underestimate the average MUTF and hence overestimate the number of MUs. A number of improvements both experimentally (e.g. Kadrie et al., 1976; Stashuk et al., 1994) and empirically (e.g. Daube, 1995; Major et al., 2007) have been proposed to try to deal with the alternation problem but, despite these improvements, each method oversimplifies the data generating mechanism and there is no gold-standard averaging approach; Bromberg (2007) and Gooch et al. (2014) provide thorough discussions on these approaches to MUNE.

A desire for a more complete model for the data generating mechanism motivated the comprehensive approach to MUNE in Ridall et al. (2006), which proposed three assumptions:

- A1** MUs fire independently of each other and of previous stimuli in an all-or-nothing response. Each MU fires precisely when the stimulus intensity exceeds a random threshold whose distribution is unique to that MU, with a sigmoidal cumulative distribution function, called an excitability curve.
- A2** The firing of a MU is characterised by a MUTF which is independent of the size of the stimulus that caused it to fire, and has a Gaussian distribution with an expectation specific to that MU and a variance common to all MUs.
- A3** The measured WMTF is the superposition of the MUTFs of those MUs that fired, together with an independent Gaussian baseline component.

From these assumptions, Ridall et al. (2006) proposed a set of similar statistical models each of which assumed a different *fixed* number of MUs. MUNE thus reduced to selection of a best model, for which the Bayesian information criterion was used. The class of methods which performs MUNE within a Bayesian framework is commonly referred to as Bayesian MUNE.

In a subsequent paper, Ridall et al. (2007) extended the method by constructing a reversible jump Markov chain Monte Carlo (RJMCMC) algorithm (Green, 1995) to sample from the MU-number posterior mass function directly. However, its implementation is highly challenging with slow and uncertain convergence particularly when the muscle has many MUs. This is partly attributed to difficulty in defining efficient and meaningful transitions between models, with transition rates found to be 0.5–2% (Andrieu, 2007). The between-model transition rate was improved in Drovandi et al. (2014) where it was noticed that under Assumption A1, for a given stimulus, the majority of MUs are either almost certain to fire or almost certain to not fire. Approximating this near certainty by absolute certainty substantially reduced the size of the sample space, permitting marginalisation in the calculation of between-model transition probabilities, and increasing the acceptance rate to 9.2% in simulated examples. Nevertheless, substantial issues over convergence remain as, for example, the parameter posterior distributions for models with more than the true number of MUs are multimodal.

In this paper, slight alterations of the neuromuscular assumptions permit the development of a fully adapted sequential Monte Carlo (SMC) filter called SMC-MUNE. As in Ridall et al. (2006), the principal inference targets are separate estimates of the marginal likelihood for a range of neuromuscular models, each assuming a different number of MUs.

The paper proceeds as follows. Section 2 presents the neuromuscular model of Ridall et al. (2006) for a fixed number of MUs and defines the priors for the model parameters. Section 3 describes the SMC-MUNE method itself. Section 4 assesses the performance of the SMC-MUNE method for 200 simulated datasets. Closer examination of cases where the point estimate of the number of MUs was incorrect revealed two classes of error; an example in each are investigated in detail. Section 4.3 presents a comparison of the SMC-MUNE method against the RJMCMC algorithm for the simulated datasets where the truth is known. Section 5 applies the SMC-MUNE method to data (collected using the method in Casella et al. (2010)) from a rat tibial muscle that has undergone stem cell therapy. Section 6 concludes the paper with a discussion on the effectiveness of SMC-MUNE and of potential avenues for improvement.

2. The neuromuscular model and prior specification

This section expands on Assumptions A1–A3 and form the model of the neuromuscular system for a given fixed number of MUs. Section 2.1 introduces the notational convention. Section 2.2 presents the neuromuscular model under the assumptions of Ridall et al. (2006), and Section 2.3 defines the prior distributions for the model parameters.

2.1. Notation

The total number of MUs operating the muscle of interest is denoted by u and a particular MU is indexed by j . An EMG dataset consists of T measurements whereby the datum for the t th test, $t = 1, \dots, T$, consists of the applied stimulus s_t and resulting WMTF y_t . The dataset is re-ordered such that the observation $y_1, \dots, y_{\tau-1}$ define baseline measurements with $s_t = 0$ for $t = 1, \dots, \tau - 1$, followed by an overall WMTF y_τ corresponding to the supramaximal stimulus $s_\tau = \max_t(s_t)$ where all u MUs are known (by the clinician) to have fired. The remaining measurements appear in order of increasing stimulus.

The reaction of MU j to stimulus s_t is $x_{j,t}$, which is 1 if MU j fires, and hence contributes to the y_t measurement, and 0 otherwise. The u -vector of indicators $\mathbf{x}_t = (x_{1,t}, \dots, x_{u,t})^\top$ defines the firing vector of the MUs in response to stimulus s_t . Given the experimental set-up, it is assumed that for each $j = 1, \dots, u$, $x_{j,t} = 0$ ($t = 1, \dots, \tau - 1$) and $x_{j,\tau} = 1$.

A sequentially indexed set of elements, vectors or scalars shall be represented using the subscript notation $a_{1:t} := \{a_1, \dots, a_t\}$. The vectors $\mathbf{0}$ and $\mathbf{1}$ have all elements zero or unity, respectively. The indicator function $\mathbb{I}_A(x)$ is 1 if $x \in A$ and 0 otherwise.

2.2. The neuromuscular model

Guided by assumptions A1–A3 of Ridall et al. (2006), our state-space neuromuscular model for the WMTF observations based on a fixed number of MUs, u , is as follows.

$$X_{j,t} | s_t, \eta_j, \lambda_j \sim \text{Bern}[F(s_t; \eta_j, \lambda_j)], \quad (1)$$

$$Y_{j,t} | \mathbf{X}_t = \mathbf{x}_t, \bar{\mu}, \bar{\nu}, \mu, \nu \sim \text{N}(\bar{\mu} + \mathbf{x}_t^\top \mu, \bar{\nu}^{-1} + \nu^{-1} \mathbf{x}_t^\top \mathbf{1}) \quad (2)$$

The WMTF in (2) is the sum of independent Gaussian contributions, firstly, from a baseline effect of $N(\bar{\mu}, \bar{\nu}^{-1})$ and, secondly, from each MU that fires. If the j th MU fires then it makes a $N(\mu_j, \nu^{-1})$ contribution to the WMTF. For crucial tractability reasons, however, we set $\bar{\nu}^{-1} = 0$ whenever at least one unit fires ($\mathbf{x}_t^\top \mathbf{1} > 0$). Neglecting the variability in the baseline noise whenever at least one MU fires is a reasonable approximation since in real experiments, because of the precautions detailed in Section 1, the variance (and expectation) of the baseline noise are generally much smaller than the variability in response from a given MU when it fires. For example, Henderson et al. (2006) finds a ratio of an order of magnitude (see also Appendix A). The parameters $\mu = (\mu_1, \dots, \mu_u)^\top$, ν , $\bar{\mu}$, $\bar{\nu}$ are collectively referred to as the *observation parameters*. In (1) the non-decreasing sigmoidal function of the stimulus, F , is called the *excitability curve* (Brown and Milner-Brown, 1976). The *excitability parameters* for the j th MU, η_j and λ_j , characterise its excitation features whereby, conditional on these values, firing events are independent between MUs. The acyclic graph in Figure 2 depicts the dependencies within the neuromuscular model. Key to the strategy in this paper is that the observational and excitability parameters are conditionally independent given the unobserved firing events $\mathbf{x}_{1:T}$.




Figure 2: Directed acyclic graph of the neuromuscular model for a fixed number of motor units, u . Arrows denote direct dependencies between known data (square nodes) and unknown parameters and states (circle nodes). Pallets indicate repeated cases according to the stated index.

The excitability curve is parameterised by its median, η and the reciprocal gradient at the median: $F(s = \eta; \eta, \lambda) = 1/2$, and $F'(s = \eta; \eta, \lambda) = 1/\lambda$. Ridall et al. (2006) specifies the excitability curve as $F(s) = \Phi[\delta(s - \eta)]$ where $\Phi(x)$ denotes the standard Gaussian cumulative distribution function (cdf). Evidence for this formulation (Hales et al., 2004) focused on the central structure of the excitability curve by applying a binned chi-squared goodness-of-fit test. Such an approach has limited ability to differentiate this formulation from alternatives as distinguishing evidence chiefly arises from tail events. Moreover, the Gaussian assumption allows a small, albeit potentially negligible, probability of a spurious firing event when no stimulus is applied. Given this contradiction with the experimental design, the following log-logistic excitability curve is used:

$$F(s; \eta, \lambda) = \left[1 + \left(\frac{s}{\eta} \right)^{-4\eta/\lambda} \right]^{-1}. \quad (3)$$

The inference method in Section 3.2 is, however, applicable for any curve.

2.3. Prior distributions

The excitability parameters of individual MUs are assumed to be independent *a priori*. For some upper limits η_{\max} and λ_{\max} , the excitability parameters are assigned vague independent beta prior distributions:

$$\frac{\eta}{\eta_{\max}} \sim \text{Beta}(1.1, 1.1), \quad \frac{\lambda}{\lambda_{\max}} \sim \text{Beta}(1.1, 1.1). \quad (4)$$

The shape parameters are chosen so that the densities are uninformative yet tail off towards the boundaries. The location upper bound is conservatively set just

greater than the supramaximal stimulus, $\eta_{\max} = 1.1s_\tau$, and the upper bound on the scale parameter of the excitability curve is set to $\lambda_{\max} = 0.04s_\tau$; see Appendix A for further discussion and Sections 4.2 and 5 for investigations of the sensitivity of inferences to changes in λ_{\max} .

The prior distributions for the observation parameters are:

$$\begin{aligned} \bar{\nu} &\sim \text{Gam}(\bar{a}_0, \bar{b}_0), & \bar{\mu}|\bar{\nu} &\sim \text{N}(\bar{m}_0, \bar{\nu}^{-1}\bar{c}_0), \\ \nu &\sim \text{Gam}(a_0, b_0), & \mu|\nu &\sim \text{MVN}_u(\mathbf{m}_0, \nu^{-1}C_0). \end{aligned} \quad (5)$$

All hyper-parameters are strictly positive scalars except for the real-valued scalar expectation \bar{m}_0 , u -vector \mathbf{m}_0 and $u \times u$ positive definite matrix C_0 . The prior distributions defined for the precision parameters are consistent with Ridall et al. (2006). However, the prior for both baseline and MUTF expectations differ from the gamma definition of Ridall et al. (2006). The tractability reasons for adopting Gaussian rather than gamma priors are detailed in Section 3.1; the problems that arise from the support now including the whole real line are addressed in Section 3.3. Because baseline observations are assimilated first, it is possible to incorporate our *a priori* knowledge that $\bar{\nu} \gg \nu$ whilst maintaining tractability, via our specification of (a_0, b_0) ; see Appendix A.

The range of MU numbers, $u = 1, \dots, u_{\max}$, defines a set of neuromuscular models. Previous Bayesian MUNE methods specified a uniform prior on the model space. However, for models that assume more MUs than there exists in truth or where there is alternation the posterior distribution is almost certain to possess multiple modes. This relates to an issue of label switching (e.g. Stephens, 2000; Aitkin, 2001) where each mode describes a variation on the same model, or even an identical model, leading to an artificial inflation of the marginal likelihood. To counteract this we introduce a preference for smaller models by specifying a $\text{Geom}(1/2)$ prior, truncated at u_{\max} : $\pi(u) \propto 2^{-u}$ for $u = 1, \dots, u_{\max}$.

3. Methodology for SMC-MUNE

The SMC-MUNE procedure detailed in this section relies upon tractable updates for the parameters given the firing vector. Tractability for the excitability parameters is based on a reasonable approximation to the ideal model defined in Section 2, and is brought about, effectively, through an approximation to the prior specification. Tractability for the firing parameters is achieved through approximation by quadrature. An overview of the methodology is provided first, with details about each part given subsequently. Adapting terminology from sequential inference, the index t shall henceforth be referred to as ‘time’.

3.1. Overview

The ultimate aim is to calculate and compare the posterior model probabilities for a range of models, each with a different number of MUs, u . Hence, for a given model with u MUs, the target for inference is its marginal likelihood,

$f(y_{1:T}|s_{1:T})$; throughout this section, for notational simplicity, we suppress the dependence on u . The marginal likelihood can be expressed as a product of sequential predictive factors with each defined by:

$$f(y_t|y_{1:t-1}, s_{1:t}) = \sum_{\mathbf{x}_{1:t-1} \in \mathcal{X}_{1:t-1}} f(y_t|\mathbf{x}_{1:t-1}, y_{1:t-1}, s_{1:t}) \mathbb{P}(\mathbf{x}_{1:t-1}|y_{1:t-1}, s_{1:t-1}) \quad (6)$$

where $\mathcal{X}_{1:t} = \{0, 1\}^{ut}$ denotes the space for the sequence of vectors of historical firing events.

The inference scheme is based upon two key points. Firstly, the observation and excitability parameters are conditionally independent given the set of firing events $\mathbf{x}_{1:T}$. Secondly, conditional on \mathbf{x}_t and with the simplifying assumption on $\bar{\nu}$ detailed in Section 2.2, the priors for the observation parameters in (5) are conjugate for the likelihood in (2), leading to standard tractable updates (see Appendix B) and tractable marginals (see (8)). Indeed, subject to the approximation, the posterior distribution for the observational parameters after assimilating $y_{1:t}$ and conditional on $\mathbf{x}_{1:t}$ is:

$$\begin{aligned} \bar{\nu}|y_{1:t}, \mathbf{x}_{1:t} &\sim \text{Gam}(\bar{a}_t, \bar{b}_t), & \bar{\mu}|\bar{\nu}, y_{1:t}, \mathbf{x}_{1:t} &\sim \text{N}(\bar{m}_t, \bar{\nu}^{-1}\bar{c}_t), \\ \nu|y_{1:t}, \mathbf{x}_{1:t} &\sim \text{Gam}(a_t, b_t), & \mu|\nu, y_{1:t}, \mathbf{x}_{1:t} &\sim \text{MVN}_u(\mathbf{m}_t, \nu^{-1}C_t), \end{aligned} \quad (7)$$

where $\bar{\mathcal{A}}_t := \{\bar{a}_t, \bar{b}_t, \bar{m}_t, \bar{c}_t\}$ and $\mathcal{A}_t := \{a_t, b_t, \mathbf{m}_t, C_t\}$.

Given (7) the marginal likelihood for the observation y_t conditional on the firing vector \mathbf{x}_t and sets $\bar{\mathcal{A}}_{t-1}$ and \mathcal{A}_{t-1} has the tractable form:

$$f(y_t|\mathbf{x}_t, \bar{\mathcal{A}}_{t-1}, \mathcal{A}_{t-1}) = \begin{cases} \mathfrak{t}\left[y_t; \bar{m}_{t-1}, \frac{\bar{b}_{t-1}}{\bar{a}_{t-1}}(\bar{c}_{t-1} + 1), 2\bar{a}_{t-1}\right] & \text{if } \mathbf{x}_t = \mathbf{0}, \\ \mathfrak{t}\left[y_t; \bar{m}_{t-1} + \mathbf{x}_t^\top \mathbf{m}_{t-1}, \frac{b_{t-1}}{a_{t-1}}q_t, 2a_{t-1}\right] & \text{otherwise,} \end{cases} \quad (8)$$

where $q_t = \mathbf{x}_t^\top \mathbf{1} + \mathbf{x}_t^\top C_{t-1} \mathbf{x}_t$. Here, $\mathfrak{t}(y; m, v, n)$ denotes the Student's t-density function on n degrees of freedom with centrality parameter m and scaling factor \sqrt{v} . The statistics $\bar{\mathcal{A}}_{t-1}$ and \mathcal{A}_{t-1} are deterministic functions of $y_{1:t-1}$ and $\mathbf{x}_{1:t-1}$, and are sufficient in that $f(y_t|\mathbf{x}_{1:t}, y_{1:t-1}) \equiv f(y_t|\mathbf{x}_t, \bar{\mathcal{A}}_{t-1}, \mathcal{A}_{t-1})$.

The posterior-predictive probability mass function for the next excitation vector, $\mathbb{P}(\mathbf{x}_t|\mathbf{x}_{1:t-1}, y_{1:t-1}, s_{1:t}) = \mathbb{P}(\mathbf{x}_t|\mathbf{x}_{1:t-1}, s_{1:t})$, is given by the following intractable marginalisation:

$$\mathbb{P}(\mathbf{x}_t|\mathbf{x}_{1:t-1}, s_{1:t}) = \int \mathbb{P}(\mathbf{x}_t|\eta_{1:u}, \lambda_{1:u}, s_t) \pi(\eta_{1:u}, \lambda_{1:u}|\mathbf{x}_{1:t-1}, s_{1:t-1}) d(\eta_{1:u}, \lambda_{1:u}) \quad (9)$$

where $\pi(\eta_{1:u}, \lambda_{1:u}|\mathbf{x}_{1:t-1}, s_{1:t-1})$ is the posterior for the excitability parameters given the firing vectors to time $t-1$. Section 3.2 presents a fast numerical quadrature scheme for evaluating (9) to any desired accuracy.

Marginalising the parameters in (8) and (9) together provide the predictive:

$$f(y_t | \mathbf{x}_{1:t-1}, y_{1:t-1}, s_{1:t}) = \sum_{\mathbf{x}_t \in \mathcal{X}_t} f(y_t | \mathbf{x}_{1:t}, y_{1:t-1}, s_{1:t}) \mathbb{P}(\mathbf{x}_t | \mathbf{x}_{1:t-1}, s_{1:t}), \quad (10)$$

Combining (10) with the historical firing event mass function $\mathbb{P}(\mathbf{x}_{1:t} | y_{1:t}, s_{1:t})$ would provide the quantity $f(y_t | y_{1:t-1}, s_{1:t})$ in (6) as required. However, it is infeasible to track the historical mass function as the dimension of the event space, $\mathcal{X}_{1:t}$, increases exponentially with time. Instead, combining (8), (9) and (10) gives the conditional mass function for the current firing vector given all previous firing vectors and all MUTFs to date:

$$\mathbb{P}(\mathbf{x}_t | y_{1:t}, \mathbf{x}_{1:t-1}, s_{1:t}) = \frac{f(y_t | \mathbf{x}_{1:t}, y_{1:t-1}, s_{1:t}) \mathbb{P}(\mathbf{x}_t | \mathbf{x}_{1:t-1}, s_{1:t})}{f(y_t | \mathbf{x}_{1:t-1}, y_{1:t-1}, s_{1:t})}. \quad (11)$$

Expressions in (10) and (11) together lead to a fully adaptive sequential Monte Carlo (SMC) sampler where the historical firing event mass function is approximated by the particle set $\{\mathbf{x}_{1:t}^{(i)}\}_{i=1}^N$, for a suitably large N , that is recursively updated for $t = 1, \dots, T$. Algorithm 1 presents the auxiliary SMC sampler (Pitt and Shephard, 1999) which, given the set of samples drawn from $\mathbf{X}_{1:t-1} | y_{1:t-1}, s_{1:t-1}$, creates an unweighted sample from the filtering distribution $\mathbf{X}_{1:t} | y_{1:t}, s_{1:t}$, and approximates (6) via Monte Carlo so as to update the marginal likelihood estimate $\hat{f}(y_{1:t} | s_{1:t})$.

Algorithm 1 Fully adapted SMC sampler for incorporating the observation y_t given that $y_{1:t-1}$ have already been assimilated.

- 1: **for** i in $1, \dots, N$ **do** ▷ Weight
 - 2: $\omega_t^{(i)} = f(y_t | \mathbf{x}_{1:t-1}^{(i)}, y_{1:t-1}, s_{1:t})$ ▷ Using (10)
 - 3: $\bar{\omega}_t^{(i)} = \omega_t^{(i)} / \sum_k \omega_t^{(k)}$.
 - 4: Sample auxiliary indices $\{\zeta^{(i)}\}_{i=1}^N$ with probabilities $\{\bar{\omega}_t^{(i)}\}_{i=1}^N$. ▷ Resample
 - 5: **for** i in $1, \dots, N$ **do** ▷ Propagate
 - 6: Sample $\mathbf{x}_t^{(i)}$ with probability $\mathbb{P}(\mathbf{x}_t | y_t, \mathbf{x}_{1:t-1}^{(\zeta_i)}, s_{1:t})$. ▷ Using (11)
 - 7: Set $\mathbf{x}_{1:t}^{(i)} = (\mathbf{x}_{1:t-1}^{(\zeta_i)}, \mathbf{x}_t^{(i)})$.
 - 8: Set $\log \hat{f}(y_{1:t} | s_{1:t}, u) = \log \hat{f}(y_{1:t-1} | s_{1:t-1}, u) - N \log \sum_i \omega_t^{(i)}$. ▷ Update \hat{f}
-

Although primary interest lies in the marginal-likelihood estimate, parameter inference is also available to assist in assessing the quality of fit. The final particle set $\{\mathbf{x}_{1:T}^{(i)}\}_{i=1}^N$ is equivalent to an N -component Gaussian-gamma mixture approximation to the posterior distribution for the observation parameters. A similar approximation to the posterior distribution for the excitability parameters is derived in Section 3.2.

3.1.1. Equivalent particle specification and degeneracy

The Bayesian conjugate structure for the observation process suggests storing and updating the sufficient statistics when assimilating the latest obser-

vations. Given the prior statistics $\bar{\mathcal{A}}_0$ and \mathcal{A}_0 , there is a deterministic map from $(\mathbf{x}_{1:t-1}, y_{1:t-1})$ to $(\bar{\mathcal{A}}_{t-1}, \mathcal{A}_{t-1})$. Hence estimates relating to the observation process at time t are equivalently expressed with respect to the samples $\{\bar{\mathcal{A}}_{t-1}^{(i)}, \mathcal{A}_{t-1}^{(i)}, \mathbf{x}_t^{(i)}\}_{i=1}^N$. The storage required for this set representation does not increase with number of observations assimilated. Since the method relies on these sufficient statistics, Algorithm 1 may be considered as a case of particle learning (Carvalho et al., 2010). For notational clarity, however, the particle set is described in terms of the historical firing events, $\mathbf{x}_{1:t-1}$, unless otherwise specified.

Assimilating the observation, y_τ , at the supramaximal stimulus, s_τ , before any of the other non-baseline observations ensures an initial update of μ for *all* MUs from the initial vague prior. After this, each MUTF location statistic is $m_j \approx y_\tau/u$, which guides predictions in (8), and hence (10), when a new MU fires. This helps to mitigate against the inevitable particle degeneracy that occurs with particle learning. Further mitigation is achieved by iteratively re-running the algorithm with more particles until inferences are stable (Appendix C).

3.2. Details for the firing vector and excitability parameters

At time $t - 1$, each particle sample consists of a historical sequence of firing events for all MUs, and from this an associated joint posterior for the firing parameters, $\eta_{1:u}$ and $\lambda_{1:u}$, is derived. A representation of the distribution for the excitability parameters is sought that should (a) permit simple calculation of the firing event predictive in (9), (b) be deterministically updatable when assimilating the current measurement, and (c) provide a concise and sufficient description for the posterior distribution.

From the independence of MU firing under Assumption A1 and the excitability parameter prior in (4), it follows that the predictive for the firing event \mathbf{x}_t in (9) factorises:

$$\mathbb{P}(\mathbf{x}_t | \mathbf{x}_{1:t-1}, s_{1:t}) = \prod_{j=1}^u \int \mathbb{P}(x_{j,t} | \eta_j, \lambda_j, s_t) \pi(\eta_j, \lambda_j | x_{j,1:t-1}, s_{1:t-1}) d(\eta_j, \lambda_j), \quad (12)$$

where the posterior for the j th MU's excitability parameters at time $t - 1$ is:

$$\pi(\eta_j, \lambda_j | x_{j,1:t-1}, s_{1:t-1}) \propto \prod_{r=1}^{t-1} \mathbb{P}(x_{j,r} | \eta_j, \lambda_j, s_r) \pi(\eta_j) \pi(\lambda_j). \quad (13)$$

Regardless of the excitability curve definition, this product of firing probabilities does not lead to a simple conjugate structure with a concise set of sufficient statistics for the posterior distribution. To address this issue, the following approximation is used:

- B** For each MU, store and update at each time point a surface proportional to the posterior density $\pi(\eta_j, \lambda_j | x_{j,1:t-1}, s_{1:t-1})$ at a set of grid of points on

a regular rectangular lattice \mathcal{G} spanning the excitability parameter space. For general (η_j, λ_j) , approximate the right-hand side of (13) using bilinear interpolation from the four nearest grid points.

Under this assumption, let $h(\eta, \lambda)$ be the right-hand side of (13); then $\tilde{h}(\eta, \lambda)$, the interpolated surface specified using points on the unit square in which (η, λ) resides, is:

$$\tilde{h}(\eta, \lambda) = (1 - \eta)(1 - \lambda)h(0, 0) + (1 - \eta)\lambda h(0, 1) + \eta(1 - \lambda)h(1, 0) + \eta\lambda h(1, 1),$$

with a similar approximation for $\mathbb{P}(x_{j,t} | \eta_j, \lambda_j, s_t) h(\eta_j, \lambda_j)$ based on interpolating this between grid points. The resulting approximations for the normalising constant in (13) and the integral in (12), therefore, correspond to iterative (over the two dimensions) application of the compound trapezium rule. This approach provides a deterministic updating procedure for maintaining the excitability posterior density up to a constant of proportionality for each point on the regular lattice.

A naïve implementation of the above scheme would evaluate the posterior density for each grid point, MU and particle sample. However, two posterior densities will only differ if the historical firing events differ. Consider any two particles, i and i' , each with an associated MU, j and j' respectively, that have identical firing histories: $x_{j,1:t}^{(i)} = x_{j',1:t}^{(i')}$. Since the priors for the excitability parameters are identical for all MUs, the posterior distribution for these two MUs on these two particles are identical. Efficiency gains are therefore achieved by storing a single grid of values for each unique firing pattern to date.

3.3. Improving the marginal likelihood estimate

The following post-processing development is motivated by the analysis of a particular simulated dataset where the point estimate for the number of MUs is one greater than the true number. The detailed analysis in Section 4.1 shows that the extra MU has a very weak expected MUTF and that it, effectively, acts simply to increase the variability in the response. The problem arises because the u -vector μ has a Gaussian prior which, to allow reasonable uncertainty across the typical range of believable MUTF contributions, also places a non-negligible prior mass at low and negative values. Negative expectations for an individual MU need not be prohibited by the data provided that the MU is always inferred to fire alongside another MU with a positive expectation of similar or larger magnitude. The fact that the parameter support permits this possibility potentially increases the marginal likelihood for a model which is larger than that necessary to explain the data.

Guaranteeing expected MUTFs greater than some minimum (Bromberg, 2003; Major et al., 2007) would require a change to the likelihood and would hence prevent the tractable updates described in Section 3.1. Within the constraints of the algorithm overviewed in Section 3.1, the natural mechanism for preventing these undesirable scenarios is via post-processing: the conditional prior for $\mu | \nu$ in (5) is re-calibrated by truncating it to the region $M = [\mu_{\min}, \infty)^u$

for some minimum MUTF $\mu_{\min} > 0$. It follows that the re-calibrated marginal prior for μ is:

$$\tilde{\pi}(\mu) = \frac{1}{\pi(M)}\pi(\mu)\mathbb{I}_M(\mu), \quad (14)$$

where $\pi(\mu)$ is the multivariate Student's t-density centred at \mathbf{m}_0 with shape matrix $\frac{b_0}{a_0}C_0$ and $2a_0$ degrees of freedom, and, with a slight abuse of notation, $\pi(M) = \int_M \pi(\mu)d\mu$. The effect on the marginal likelihood from the prior re-calibration is examined in Proposition 1, which is proved in Appendix D.

Proposition 1. *Let $f(y_{1:T}|s_{1:T})$ denote the marginal likelihood defined in Section 3.1. The re-calibrated marginal likelihood, denoted by $\tilde{f}(y_{1:T}|s_{1:T})$, resulting from truncating the prior for μ in (5) to region M in specifying (14) is:*

$$\tilde{f}(y_{1:T}|s_{1:T}) = \frac{\pi(M|y_{1:T}, s_{1:T})}{\pi(M)}f(y_{1:T}|s_{1:T}) = \frac{\pi(M|y_{1:T}, s_{1:T})}{\pi(M|y_{1:\tau-1}, s_{1:\tau-1})}f(y_{1:T}|s_{1:T}), \quad (15)$$

where $\pi(M|y_{1:t}, s_{1:t}) = \int_M \pi(\mu|y_{1:t}, s_{1:t})d\mu$.

$\pi(M|y_{1:\tau-1}, s_{1:\tau-1})$ is an orthant probability for the multivariate Student-t distribution. In contrast, the numerator in (15) is estimated from the N -component-mixture approximation of the posterior distribution by the final particle set: $\hat{\pi}(M|y_{1:T}, s_{1:T}) = \frac{1}{N} \sum_{i=1}^N \pi(M|\mathbf{x}_{1:T}^{(i)}, y_{1:T})$.

4. Simulation study

The performance of the SMC-MUNE algorithm is now assessed using 200 simulated datasets, 20 for each true number of MUs of $u^* = 1, \dots, 10$. Each dataset consists of $T = 220$ measurements with $\tau = 21$ so that the first 20 observations correspond to the baseline, $s_t = 0$ V, and these are followed by the supramaximal stimulus $s_{21} = 40$ V.

All MUs are excited according to the log-logistic curve in (3) with MU parameters simulated anew for each dataset as follows: $\eta_j \sim \text{Unif}(5, 40)$, $\lambda_j \sim \text{Gam}(2, 8)\mathbb{I}(\lambda_j < 10)$, $\mu_j \sim \text{N}(40, 20^2)\mathbb{I}(\mu_j > 20)$, $\nu^{-1} \sim \text{Unif}(1, 5)$. The measurement units for excitation parameters are all in V and the expected MUTFs are in mN with variance parameter in mN². Parameters were independent except for the following constraints, where (j) is the index of the MU with the j th highest η value: $\eta_{(j)} - \eta_{(j-1)} > 2$ (neighbours must be separate) and $|\mu_{(j)} - \mu_{(j-1)}| > 4$ (neighbours have distinct expectations). To test the ability of SMC-MUNE, these values ensure a greater range in MUTF contributions and excitation than is typically observed in practice. For example, the average coefficient of variation is 11.2% (as opposed to 1.65% (Hales et al., 2004)). This resulted in a high proportion of datasets, 79%, containing at least one alternation event. Additional noise was generated as in (2) with $\bar{\mu} = 0$ mN and $\bar{\nu}^{-1} = 0.25^2$ mN².

To each dataset, a set of neuromuscular models was fitted with a number of MUs, u , ranging from 1 to $u_{\max} = 12$. The sufficient statistics for the parameter prior distributions are provided in Appendix A.1. Control on the Monte Carlo variability and numerical integration error was ensured by iteratively increasing the particle set size, N , and the lattice size until estimates of the marginal likelihood were stable; see Appendix C for further details. The MU number marginal posterior is summarised by the maximum *a posteriori* (MAP) point estimate, $\hat{u} = \operatorname{argmax}_u \hat{\mathbb{P}}(U|y_{1:T}, s_{1:T})$, and the 95% highest posterior credible set (HPCS) to quantify estimation uncertainty; the minimal set of models where their total posterior probability is at least 95%. In addition, the estimated posterior probability for the true model, $p_{u^*} = \hat{\mathbb{P}}(U = u^*|y_{1:T}, s_{1:T})$, is evaluated.

Table 1 presents summaries of the mass functions of the number of MUs and descriptions of the resource required as functions of the true number of MUs. The MAP estimate corresponded to the true number of MUs for all datasets generated from five or fewer MUs, and for most of these datasets the HPCS contained only the true model. For true sizes of greater than five the MAP estimate was correct for 81 of the 100 datasets and the HPCS contained the truth for all but two datasets.

It is unsurprising that the uncertainty in the MU-number increases with the true number of MUs; this is visible both as an increase in the average size of the HPCS and a reduction in the average posterior probability for the true number. In addition, both the number of particles required to control Monte Carlo variability and the size of the numerical lattice required for accurate numerical integration also increase with the true number of MUs. This demonstrates the challenge of MUNE for large neuromuscular systems.

Sensitivity to the prior distribution on the number of MUs was investigated by, replacing the geometric prior with the flat prior used in previous Bayesian MUNE methods. Posterior estimates for the number of MUs were mostly the same, with an additional 4 datasets where the MAP estimates correctly identified the truth. Closer assessment showed that in these cases, under the geometric prior, the estimated posterior probability for the truth was only slightly less than that of the MAP estimate.

Of the 19 datasets from Table 1 where the MAP estimate \hat{u} did not correspond to the truth, u^* , one dataset had $\hat{u} > u^*$ with the rest (including the two outliers) having $\hat{u} < u^*$. The stimulus-response curves for the first case and a typical example of the second are presented in Figure 3 and are discussed in turn below.

4.1. Overestimation

The first dataset, D1, contains $u^* = 7$ MUs but the SMC-MUNE MAP estimate is $\hat{u} = 8$. The posterior probability of the true model is 14.9% and this model, along with the larger 9 MU model, is contained with a 95% HPCS. Parameter estimates for the MAP model (Table 2) show that the penultimate MU has a median expected twitch force of 9.6 mN with a credible upper bound of 15.7 mN, much lower than the 20 mN simulation threshold. Figure 4 presents

Table 1: Summary of the MU-number posterior mass functions and required numerical resource for 200 simulated datasets.

Number of MUs, u^*	≤ 5	6	7	8	9	10
No. where $\hat{u} = u^*$	100	19	19	16	15	12
No. where u^* in HPCS	100	20	20	19	19	20
Avg. size of HPCS	1.11	1.70	2.10	2.05	2.35	2.45
Avg. \hat{p}_{u^*} (%)	97.95	89.20	80.45	69.42	62.70	54.68
Avg. particle set size for u^*	5000	5250	6000	7500	7500	8250
Avg. particle set size for \hat{u}	5000	5000	6000	7500	9250	7750
Avg. $n \times n$ lattice size for u^*	30.0	30.5	30.5	32.0	32.5	32.0
Avg. $n \times n$ lattice size for \hat{u}	30.0	30.0	30.5	32.0	35.0	32.0

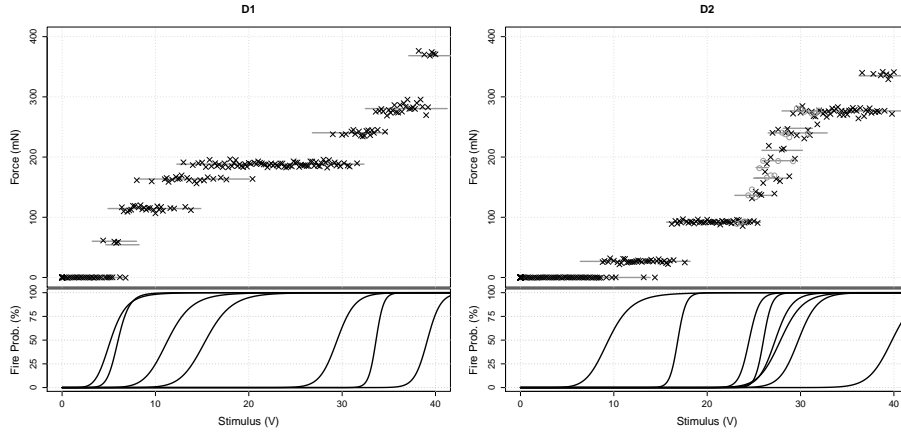


Figure 3: Stimulus-response curve (top) for the simulated data with lines representing the expected WMTF over the stimuli intervals where the joint firing probability is greater than 5% according to the individual excitability curves (bottom). Left: Dataset D1 contains $u^* = 7$ MUs but $\hat{u} = 8$. Right: Dataset D2 contains $u^* = 8$ MUs but $\hat{u} = 7$. Circle points, additional 23 simulations over the 23–32 V alternation period involving 5 MUs.

Table 2: Expected MUTF median and 95% credible interval estimates for MUs with high excitation threshold from the true ($u^* = 7$) and MAP ($\hat{u} = 8$) models, with and without post-process truncation ($\mu_{\min} = 15$ mN) on dataset D1.

Parameter	μ_6	μ_7	μ_8	ν^{-1}
True	40.2 mN	87.9 mN	–	4.54 mN ²
$u = 7$	40.5 (37.7, 43.5)	91.2 (86.6, 95.9)	–	3.85 (3.13, 4.81)
$u = 8$	36.3 (32.4, 40.2)	9.6 (4.7, 15.7)	86.7 (80.3, 92.7)	3.22 (2.57, 4.18)
$u = 7$ & $\mu_{\min} = 15$	40.5 (37.8, 40.2)	91.3 (86.8, 95.7)	–	3.90 (3.14, 4.92)
$u = 8$ & $\mu_{\min} = 15$	35.7 (31.1, 40.0)	15.7 (15.0, 20.7)	80.4 (71.4, 86.6)	3.23 (2.60, 4.09)

the the construction of the predictive WMTF density for the true and MAP models at a 37 V stimulus. The local modes in the model containing the true MU-number correspond uniquely to particular firing combinations. In contrast, the weak MU in the MAP model principally serves to increase the variability around a specific WMTF response level rather than describing a distinct MU.

In light of these concerns, the marginal likelihood estimates are adjusted according to Section 3.3 with a conservative lower bound of $\mu_{\min} = 15$ mN to guard against small MUs that, when firing, are indistinguishable from other combinations. The corrected posterior mass function places 89.3% of the mass on the correct, seven-MU model, with 10.7% mass on the eight-MU model. The estimates of expected MUTF in Table 2 for the seven-MU model are similar to those prior to the adjustment and are still close to the true values from which the data was generated. Unsurprisingly, the prior adjustment for the wrong, eight-MU hypothesis has a significant effect on the penultimate MU and, so as to preserve the overall maximum WMTF, a small reduction in the estimated μ for its neighbouring MUs.

The prior adjustment on the other 199 datasets was determined and the effect on the quantities reported in Table 1 are positively effected (see Appendix E for details). In summary, the truth corresponded to the MAP on 83 occasions instead of 81, the averaged HPCS width is reduced and the average posterior probability for the truth is increased. Also, the 95% HPCS excluded the truth on six occasions, rather than two previously, where in each case the number of MUs was underestimated; a phenomenon which we investigate next.

4.2. Underestimation

The second dataset, D2, contains $u^* = 8$ MUs and exhibits a period of alternation between 23–32 V which involves five MUs. SMC-MUNE, however, estimates $\hat{u} = 7$ and with a posterior probability of 97.1% after applying the post-process adjustment at $\mu_{\min} = 15$ mN (95.8% before applying the adjustment). The main source for underestimation of u is the overestimation of the excitability scale parameter for the fourth MU (λ_4 : 1.8, $\hat{\lambda}_4$: 4.5 (1.8, 7.6)).




Figure 4: Predictive density (thick line) at stimulus 37 V from the seven (left) and eight (right) MU model without post-process adjustment. Thin lines identify the contribution to the predictive for the indicated firing combinations associated to the final few MUs. In both cases, the first five MUs fire with near certainty. Most firing combinations with negligible predictive probabilities are omitted from the plot.

Under this estimate the stimulus interval where this MU fires in a probabilistic behaviour is nearly three times wider than it should be. As a consequence, this incorrectly estimated MU acts as a surrogate for other units that have similar twitch force properties.

One potential solution is to reduce the upper bound for the scale parameter λ_{\max} in (4), thus forbidding shallow excitability curves. This, however, does not solve the problem of incorrect MUNE; see Appendix F for details.

The original analysis mis-estimated λ_4 because of the limited information available in the observations to adequately describe the period of alternation between 23–32 V which involves five MUs. To show that this is the case, an additional 23 observations were generated evenly over this interval; see Figure 3. This modest addendum is sufficient for the true model to be identified, $\hat{u} = 8$, with a posterior probability of 71.3% and better scale parameter estimates. However, the increase in computational resource required to obtain the same degree of Monte Carlo and numerical accuracy was substantial: from 5000 to 25000 particles and from a 30×30 to 50×50 lattice for the eight-MU hypothesis.

4.3. Comparison with RJMCMC

To benchmark SMC-MUNE we applied the RJMCMC method to the 200 simulated datasets. Because the model and prior in Section 2 differ in a number of subtle aspects compared to Ridall et al. (2007) and Drovandi et al. (2014), a new RJMCMC algorithm (with analogous reversible jump moves) was created to target the same posterior distribution as SMC-MUNE (see Appendix G for further details). For each dataset, three independent chains were run and a permutation test-based diagnostic that used the u values from the chains was

used to indicate convergence. Evidence of failed convergence was present for 23% of the datasets, and further diagnostics showed that at least one of the chains was stuck in a minor mode. This emphasises the known convergence and mixing problems associated with the RJMCMC procedure for MUNE (Andrieu, 2007).

From the 154 datasets where the evidence convergence diagnostic for the RJMCMC is satisfactory, the maximum posterior estimate for the number of MUs agreed with the SMC-MUNE method in 142 cases. Within this, both methods correctly identify the true MU number for 139 datasets whereas in the remaining 3 datasets the methods underestimate the true count by one (see Section 4.2). In the remaining 12 datasets where the SMC-MUNE and RJMCMC MAP estimates did not agree, yet the permutation test based on the u values from the RJMCMC chains suggested convergence, further investigation of the parameter values of the RJMCMC chains provided evidence of questionable convergence and mixing.

5. Case study: rat tibial muscle

The data for this case study was collected according to the experimental design in Casella et al. (2010): a rat tibial muscle (medial gastrocnemius) receives stem cell therapy to encourage neuromuscular activation after simulating paralysis. The two datasets, presented in Figure 1 were generated by applying stimuli for different durations. The first, R10, used 10 μ sec-duration stimuli and consist of $T = 304$ observations, including 11 baseline measurements and a maximal stimulus of 100 V. The second, R50, used 50 μ sec-duration stimuli and consists of $T = 669$ observations, including 7 baseline measurements, and with a maximal stimulus of 60 V. Since both datasets are collected from the same neuromuscular system it is expected that MUNE for R10 and R50 should be similar.

Naïve assessment of the stimulus-response curves by counting the number of distinct levels in twitch force suggests nine or ten MUs, but this ignores any potential features arising from alternation. The histogram inserts in Figure 1 present frequency in absolute difference between consecutive twitch forces when ordered by stimulus intensity. The highest frequency occurs at low differences and represents the within-MUTF variability whereas the less-frequent, larger differences appear due to the firing of different combinations of MUs.

The SMC-MUNE procedure was applied up to a maximum model size of $u_{\max} = 12$ with prior sufficient statistics given in Appendix A.1 and algorithmic parameters as specified in Appendix C. Without the post-processing correction of Section 3.3, a 9 or 10 unit model is favoured for R10 ($\mathbb{P}(U \in \{9, 10\}) \approx 0.782$) whilst a 10-unit model is favoured for R50 ($\mathbb{P}(U = 10) \approx 1.00$); however, in all of these cases spurious units with small or negative contributions to the twitch force are identified. After post-processing with $\mu_{\min} = 15$ mN there is strong agreement between the inferences from the two datasets. For both datasets, the estimated motor unit number posterior mass function (Table 3) identifies the MAP estimate as $\hat{u} = 8$, with this being the only member of the HPCSs. There is a noticeable difference in the computational resources required as the

Table 3: Posterior summary from the SMC-MUNE procedure for the rat tibial muscle using $10\mu\text{sec}$ and $50\mu\text{sec}$ duration stimuli.

dataset	R10				R50			
MUs (u)	7	8	9	10	7	8	9	10
p_u (%)	0.04	99.95	0.01	0.00	0.00	100.00	0.00	0.00
Grid ($n \times n$)	30	30	30	30	100	90	50	90
Particles (,000s)	20	5	15	5	155	100	65	115




Figure 5: Estimated excitability curves from the eight MU hypotheses for datasets R10 (left) and R50 (centre) with corresponding expected MUTF mean estimates.

MAP model for R50 required twenty times more particles and three times finer lattice than that for dataset R10. This reflects the relative sizes of the datasets, but may also relate to the relative complexities of the state-spaces for the firing vectors.

Figure 5 presents the estimated excitability curves for each of the MUs, with MUs labelled in order of increasing $\mathbb{E}[\eta|y_{1:T}]$. The location parameters under the $50\mu\text{sec}$ duration stimuli are approximately one quarter of the corresponding parameters under $10\mu\text{sec}$ duration stimuli. This corresponds to Weiss’s law (Bostock, 1983), which relates the excitation of the neuron to the charge built-up in the cell. Despite this, it is clear that all MUs except the last are excited within a low and narrow stimulus window, consistent with the shape of Figure 1.

Comparisons across the two datasets of the coefficients of variation of the excitability curves and the most probable firing patterns at each visibly distinct response level in Figure 1 are given and discussed in Appendix H.

6. Discussion

This paper presents a new sequential Bayesian procedure for motor unit number estimation (MUNE), the assessment of the number of the operating

motor units (MUs) from an electromyography investigation into muscle function. The fully adapted sequential Monte Carlo (SMC) filter uses the approximate conditional conjugacy of the twitch process. SMC-MUNE estimates the marginal likelihood for the neuromuscular model based on a fixed number of MUs. From this, MUNE is then performed by comparing the evidence between competing MU-number hypotheses. As is demonstrated in Sections 4.1 and 4.2, SMC-MUNE also allows detailed scrutiny of the quality of each model fit.

SMC-MUNE performed well on simulated data; however it is possible for the expected twitch force for one or more MUs to be estimated as negligible or negative, resulting in overestimation of the number of MUs. This led to the development of a post-process correction that restricts the parameter space. Alternatively, difficulty in estimating the underlying process during a period of alternation (where the same stimulus, applied repeatedly can lead to several different combinations of MUs firing) was seen to lead to underestimation. This issue was resolved when additional data points were sampled from the region of alternation. For the majority of the simulated datasets the posterior estimates from the SMC-MCMC agreed with the estimates obtained from an RJMCMC algorithm similar to those in (Ridall et al., 2007; Drovandi et al., 2014). However, the RJMCMC algorithm showed evidence of failed convergence for about one quarter of the datasets.

Independent application of SMC-MUNE to two datasets (with data collected as in Casella et al. (2010)) from the same neuromuscular system resulted in the same estimate for the number of MUs. The models fits were consistent, although closer examination identified minor variations in parameter estimates and firing patterns, reflecting subtle, known differences between the two datasets.

Unlike a generic SMC algorithm, the proposed algorithm utilises the conjugate structure within the model. A naïve SMC would use a particle representation of parameters rather than of the sufficient statistics via the firing vectors. Each particle would hold a value for $3u + 3$ continuous parameters rather than a single firing history of binary values, leading to a huge increase in the size of the parameter space. A very large number of particles would be necessary to sufficiently cover this space and capture the structure of the posterior distribution.

The examples investigated in this paper involve neuromuscular systems with relatively small numbers of MUs, in which there is considerable interest in practice for some animals or in patients with amyotrophic lateral sclerosis who have limited motor function. Large and healthy muscle groups can contain hundreds of MUs (Gooch et al., 2014), yet the computational cost of SMC-MUNE increases exponentially with the assumed number of units. The computational demand arises from the evaluation of the predictive mass function for sampling the firing vectors, and to marginalise over this event space for calculating the resampling weights. Extension of SMC-MUNE to such systems might be addressed by approximating very low or high excitation probabilities by their respective certainties as in Drovandi et al. (2014). Alternatively, the excitability curve for SMC-MUNE is specified in generic terms and computational savings are possible by defining a function that has finite support.

Further computational resource is expended in numerical quadrature via the trapezium rule using a sufficiently fine numerical lattice over the excitability parameter space. A higher-order Newton-Cotes numerical integration method would produce a more accurate estimate of (12) with a reduced computational expense, but the associated interpolated density surface of piecewise polynomials would not be guaranteed to be non-negative, making an inspection of parameter estimates for assessing model fit problematic. Alternatively, quadrature on adaptive sparse grids (Bungartz and Dirnstorfer, 2003), where the grid is finer at regions of high curvature, could improve estimator accuracy over the static regular rectangular lattice. However, this would be achieved at the expense of additional implementation complexity and further approximation error when estimating the surface at in-filled lattice points.

In Section 2.2 we adjusted the observation process so that conditional on the firing vector the updates to the observation parameters were tractable. With the same conditioning, the excitability parameters could be updated using two-dimensional numerical integration. One might ask whether, in these days of fast computing, tractability and near-tractability is essential, or whether an off-the-shelf algorithm such as Hamiltonian Monte Carlo (HMC) (e.g. Neal, 2011) could be employed instead on a vastly extended state-space. The relatively slow speed and poor mixing of the algorithm of Ridall et al. (2007), which does take advantage of some tractable updates, provides a partial rebuttal, as does the fact that the state-space for the firing vectors is both discrete (so that HMC cannot be applied) and extremely large (size = 2^{uT}). More fundamentally, when the difficulty of a problem is exponential in some quantity (here u) then off-the-shelf, brute-force approach is bound to break down quickly.

The sequential aspect of the proposed methodology provides the opportunity for real-time inference that has the potential to provide in-lab assistance during experimentation. In this framework, an interim SMC-MUNE analysis could help in identifying the choice of stimulus to apply in order to collect the best evidence to distinguish between competing hypotheses, as in Section 4.2. The limitations of the present SMC-MUNE procedure to become a wholly online algorithm are the computational aspects discussed earlier and the post-processing stage to correct for potentially negligible estimates of the expected MU twitch forces. Solutions to these outstanding problems would increase the efficiency and accuracy of SMC-MUNE and, hence, the range of application.

Code for SMC-MUNE is available at <https://github.com/taylors2/SMCMUNE>.

Acknowledgements

The authors would like to thank Dr. Christine Thomas from the Miami Project to Cure Paralysis who provided the data analysed in this paper and for specialist discussions. The authors declare that there are no competing interests.

Appendix A. Further details on the prior distributions

Evidence for specifying the upper bound λ_{\max} is taken from Hales et al. (2004) where, for a Gaussian excitability curve, the coefficient of variation of a random variable whose cdf is given by the excitability curve was estimated to be 1.65%. With the log-logistic curve this corresponds to $\lambda/\eta \approx 3.64\%$. Given that $\eta \leq \eta_{\max} = 1.1s_\tau$, we deduce that $\lambda \leq 0.04s_\tau$. The limitations of the study of Hales et al. (2004), commented on by Major et al. (2007), indicate that a larger bound may be required than initially suggested, which was why sensitivity of MUNE to λ_{\max} was investigated in Sections 4.2 and 5.

After assimilating the baseline observations $y_1, \dots, y_{\tau-1}$, both $\bar{\nu}^{-1}$ and $\bar{\mu}$ are known (and known to be small) with considerable certainty. Thus, given that we also know $\bar{\nu} \gg \nu$, approximating $\bar{\nu}^{-1}$ as 0 and considering $\bar{\mu}$ to be a point mass at \bar{m} is reasonable. Furthermore, the prior for ν does not need to be set until just before the observation y_τ is assimilated. Given the tight posterior for $\bar{\nu}$ at this juncture it is, therefore, possible to incorporate the knowledge that $\bar{\nu} \gg \nu$ into the vague prior for ν (which is conceptually similar to specifying an initial joint prior on $\bar{\nu}$ and ν). Letting $\bar{\nu}_{\tau-1}^{\text{med}}$ denote the posterior median of $\bar{\nu}$ at time $\tau - 1$, tuning parameters $\epsilon \lll 1$ and $\delta \lll 1$ are chosen such that $\mathbb{P}(\nu > \epsilon\bar{\nu}) \approx \delta$ is desired. Given that

$$1 - \delta = \mathbb{P}(\nu \leq \epsilon\bar{\nu}) \approx \mathbb{P}(\nu \leq \epsilon\bar{\nu}_{\tau-1}^{\text{med}}) = \text{Gam}(b_{\tau-1}\epsilon\bar{\nu}_{\tau-1}^{\text{med}}; a_{\tau-1}), \quad (\text{A.1})$$

where $\text{Gam}(x; \alpha)$ is the cdf of a $\text{Gam}(\alpha, 1)$ variable evaluated at x . In practice we specify the prior for ν by defining a small $a_{\tau-1} = a_0$ and then solving (A.1) for $b_{\tau-1}$.

Appendix A.1. Prior hyperparameters

The prior sufficient statistics for the simulation and case studies are: $\bar{m}_0 = 0$, $\bar{c}_0 = 10^3$, $\bar{a}_0 = 0.5$, $\bar{b}_0 = 0.1$, $\mathbf{m}_0 = 40\mathbf{1}_u$, $C_0 = 10^4 I_u$ and $a_0 = 0.5$ where $\mathbf{1}_u$ is a unit u -vector and I_u is the $u \times u$ identity matrix. The statistic b_0 is defined according to (A.1) with $\delta = 0.05$ and $\epsilon = 0.2$. The upper bounds for the excitability parameter space are $\eta_{\max} = 1.1s_\tau$ and $\lambda_{\max} = 14$. In the case study, the upper bound for the scale parameter was reduced to $\lambda_{\max} = 7$.

Appendix B. Tractable updates for the observation parameters

Consider the observation model from (2). At time $t \leq \tau - 1$, when no MUs fire, $\mathbf{x}_t = \mathbf{0}$, the observation, y_t , provides no new information about μ and ν ; hence $\mathcal{A}_t = \mathcal{A}_{t-1}$ and $Y_{j,t}|\mathbf{x}_t = 0, \bar{\mu}, \bar{\nu}, \mu, \nu \sim \text{N}(\bar{\mu}, \bar{\nu}^{-1})$.

From (7), $\bar{\mu}|\bar{\nu}, y_{1:t}, \mathbf{x}_{1:t} \sim \text{N}(\bar{m}_t, \bar{\nu}^{-1}\bar{c}_t)$. Whenever at least one unit fires ($\mathbf{x}_t \neq \mathbf{0}$), approximating $\bar{\nu}^{-1}$ by 0 is, therefore, equivalent to setting $\bar{\mu} = \bar{m}$ and replacing (2) with:

$$Y_t | \bar{\nu}, \mu, \nu, \bar{\mathcal{A}}_{t-1}, \mathcal{A}_{t-1} \sim \text{N}(\bar{m}_{t-1} + \mathbf{x}_t^\top \mu, \nu^{-1} \mathbf{x}_t^\top \mathbf{1}).$$

Standard conjugate updates may, therefore, be applied for all observation parameters as follows:

$$\bar{m}_t = \frac{y_t + \bar{m}_{t-1}\bar{c}_{t-1}}{1 + \bar{c}_{t-1}}, \quad \bar{c}_t = \frac{\bar{c}_{t-1}}{1 + \bar{c}_{t-1}}, \quad \bar{a}_t = \bar{a}_{t-1} + \frac{1}{2}, \quad \bar{b}_t = \bar{b}_t + \frac{(y_t - \bar{m}_{t-1})^2}{2(1 + \bar{c}_{t-1})},$$

and

$$\begin{aligned} \mathbf{m}_t &= \mathbf{m}_{t-1} + q_t^{-1} C_{t-1} \mathbf{x}_t (y_t - \tilde{m}_t), & C_t &= C_{t-1} - q_t^{-1} C_{t-1} \mathbf{x}_t \mathbf{x}_t^\top C_{t-1}, \\ a_t &= a_{t-1} + \frac{1}{2}, & b_t &= b_{t-1} + \frac{1}{2q_t} (y_t - \tilde{m}_t)^2, \end{aligned}$$

where $q_t = \mathbf{x}_t^\top \mathbf{1} + \mathbf{x}_t^\top C_{t-1} \mathbf{x}_t$ and $\tilde{m}_t = \bar{m}_t + \mathbf{x}_t^\top \mathbf{m}_{t-1}$. Our approximation decouples the learning about the observational parameters: when no MU fires then $(\bar{\mu}, \bar{\nu})$ is updated, else (μ, ν) is updated.

Appendix C. Algorithm details

Resampling in Algorithm 1, is performed by systematic sampling on the residuals of particle weights (Hol et al., 2006). The number of particle samples and the number of rectangular lattice cells is initially $N = 5000$ and $|\mathcal{G}| = 30 \times 30$ respectively.

Accuracy in MU-number posterior mass function is managed by ensuring that for each model, u , the range in marginal log-likelihood estimate from 3 independent runs of the SMC scheme is less than 1 whenever the posterior probability is greater than 1%. If not, then the particle set is increased in steps of 5000 samples to reduce Monte Carlo variability. Once this criterion is satisfied, the lattice for the numerical integration is made finer by 10 vertices in both dimensions and the stability of the estimates to increasing grid size is checked. Instability leads to a further check of the Monte Carlo variability and, if necessary, an increase in the particle set size, and then a further increase in the number of vertices. Iteration between these two steps continues until the results are numerically stable and have a low variance. Once the minimal number of particles and grid size required for stability have been ascertained, a further 7 runs are performed using these settings and the final marginal log-likelihood estimate is then averaged from across the 10 runs.

Appendix D. Recalibrating the marginal likelihood estimate

This section contains the proof of Proposition 1 and a discussion of an alternative *online* recalibration.

Proof. Expressing the re-calibrated marginal likelihood as a marginalisation of

μ and substituting the definition (14) produces the first equality by:

$$\begin{aligned} \tilde{f}(y_{1:T}|s_{1:T}) &= \int \tilde{\pi}(\mu) f(y_{1:T}|\mu, s_{1:T}) d\mu = \frac{\int_M \pi(\mu) f(y_{1:T}|\mu, s_{1:T}) d\mu}{\pi(M)} \\ &= \frac{\pi(M|y_{1:T}, s_{1:T})}{\pi(M)} f(y_{1:T}|s_{1:T}). \end{aligned}$$

The second equality in (15) arises as the first $\tau - 1$ observations relate exclusively to the baseline. \square

There is no theoretical argument against assuming (14) from the outset. Indeed, propagating of the firing events in Algorithm 1 would account for the truncation to parameter space for μ , directing particle to a more appropriate parameter posterior approximation. However, implementing this scheme requires up to $N2^u$ orthant evaluations of the multivariate Student’s t-distribution per time step in calculating the re-sampling weights. Standard procedures for evaluating these orthant probabilities (Genz and Bretz, 2009) are expensive, so the computational time of the resulting SMC-MUNE algorithm would increase substantially.

Appendix E. Overestimation of the number of MUs

Table E.4 corresponds to the simulation study of Section 4 and presents the quantities reported in Table 1, but after the post-processing correction described in Section 4.1.

Table E.4: Summary of the MU-number posterior mass functions for 200 simulated datasets with post-process adjustment.

Number of MUs, u^*	≤ 5	6	7	8	9	10
No. where $\hat{u} = u^*$	100	19	20	17	14	13
No. where u^* in HPCS	100	20	20	18	19	17
Avg. size of HPCS	1.02	1.05	1.15	1.30	1.45	1.60
Avg. \hat{p}_{u^*} (%)	99.76	97.06	97.95	81.80	71.96	58.18

Appendix F. Underestimation of the number of MUs

SMC-MUNE underestimated the number of motor units in simulated dataset D2 (see Section 4.2). This section details an investigation into the effect of altering the excitability scale parameter λ , with MU number posterior probability estimates presented in Table F.5. Under the reduced bound the 8 MU model becomes a member of the HPCS, but the MAP estimate remains at $\hat{u} = 7$ with a high posterior probability of 94.6%. Although a further reduction to λ_{\max} might be appealing, this action is likely to be detrimental in determining good model fits. For example, the scale parameter of the first MU, with true value of 5.0 V,

is well estimated at 4.7 V whether λ_{\max} is 7 or 14. The excitability of this MU is well separated from the other curves and so a further reduction in λ_{\max} is likely to adversely affect inference away from the truth through increased uncertainty in other parameters or the introduction of spurious MUs.

Table F.5: Motor unit posterior probabilities for D2 with scale upper bound at $\lambda_{\max} = 14$ V and $\lambda_{\max} = 7$ V, and with 23 additional measurements at $\lambda_{\max} = 14$ V.

MUs (u)	$\lambda_{\max} = 14$	$\lambda_{\max} = 7$	Extra
7	96.7%	94.6%	28.7%
8	3.3%	5.4%	71.3%

Appendix G. RJMCMC algorithm for MUNE

This section outlines the particulars of the RJMCMC algorithm applied in Section 4.3. This algorithm is subtly different to Ridall et al. (2007) and Drovandi et al. (2014) due in principal to different targeted posterior distributions arising from the minor differences in the statistical model and prior specification in Section 2. The minimum expected MUTF constraint, μ_{\min} , is applied to the model from the outset and estimates are compared against the post-process corrected posteriors according to Section 3.3.

The core of the RJMCMC algorithm is based on Drovandi et al. (2014) where focus is given to the binary firing vectors, as in Section 3.1, rather than on the auxiliary excitation threshold in Ridall et al. (2007). This approach achieves better mixing properties. The RJMCMC algorithm requires that the MUs be ordered such that the excitability location parameters are increasing. This additional requirement is applied to the joint prior distribution in (4).

The details for sampling the parameters for a within-model update where the number of MUs is fixed are as follows. A draw for each of the firing vectors is sampled directly from the full conditional $x_t | \dots \sim f(y_t | \mathbf{x}_t, \bar{\mu}, \bar{\nu}, \mu, \nu) \mathbb{P}(\mathbf{x}_t, \eta_{1:u}, \lambda_1 : u, s_t)$ for $t > \tau$, whereas the vectors for for the initial baseline events, $\mathbf{x}_t = \mathbf{0}$ for $t < \tau$, and for the supramaximal event, $\mathbf{x}_\tau = \mathbf{1}$, and fixed. Samples for the baseline mean, $\bar{\mu}$, and each of the expected MUTFs, μ_1, \dots, μ_u , are drawn directly from the tractable (truncated-)Gaussian full conditional distributions. Updates to all excitability parameters, the baseline variance and within MU variance are drawn from the full conditionals using a slice sampler (Neal, 2003).

Proposals for jumping between models by adding or removing MUs is identical to Drovandi et al. (2014). A MU selected at random with weights related to the span of the excitability curve such that MUs within a period of alternation are more likely to be chosen. When increasing the number of MUs, the selected MU is split into two where one adopts the current excitability scale parameter from the original unit whilst the other is assigned a new value drawn from the prior. The excitability location and expected MUTF parameters for the new MUs are evaluated according to Ridall et al. (2007). Note that Drovandi et al. (2014) introduces an extra term into the MU selection weights whereby

only those units with expected MUTF greater than $2\mu_{\min}$ can be selected such that the resulting units both satisfy the minimum twitch size constraint. The baseline parameters, the within MU variance parameter and the parameters associated with the other MUs are unaltered.

The proposal under reducing the number of MUs is determined by the reverse mapping to that described above based on combining the selected MU and the MU with the next largest excitability location parameter. To minimise tuning, the alternative 2 \leftrightarrow 3 trans-dimensional proposal Ridall et al. (2007) was not applied. At each iteration of the algorithm a single update of the model parameters under a fixed MU number was performed, which was then followed by a single trans-dimensional proposal where the split and combine moves were selected with equal probability.

The RJMCMC algorithm was performed three times for each dataset with different initial values drawn from the prior. Each chain was executed for at least 10 times the cpu time required for the SMC-MUNE method. The first half the chains was removed as burn-in and the sum of absolute differences in the estimated number MUs posterior between the three chains were determined: $\frac{1}{3} \sum_{c=1}^3 \sum_{u=1}^{u_{\max}} |\pi_c(u|y) - \bar{\pi}(u|y)|$ where $\bar{\pi}(u|y) = \frac{1}{3} \sum_{c=1}^3 \pi_c(u|y)$. A permutation test was performed as diagnostic on chain convergence.

Appendix H. Coefficients of variation and most probable firing patterns for the rat data

To compare MUs between datasets, the coefficient of variation for the random variable associated with each excitability curve is presented in Figure H.6. Apart from the first MU, the 95% credible intervals from each dataset for a given MU overlap, suggesting similar coefficients of variation for the MUs, as might be anticipated since measurements are taken from the same neuromuscular system. These similar estimated coefficients are larger than the estimate in Hales et al. (2004), which is presented for comparison. This reflects the experimentation, where the developed neurons are less stable and are yet to restore full and healthy motor function.

Table H.6 presents the most probable firing combinations for each visibly distinct response level in Figure 1. The estimated firing behaviour of each MU, after label-swapping similarly excitable MUs for R50, are very similar between the two datasets. It can then be suggested that the level at approximately 120 mN in both datasets and at about 70 mN in R50 are potential consequences of alternation as MUs that fired in contributing to weaker WMTFs are latent in forming these response levels. However, a difference in estimated firing behaviour occurs at the 120 mN response level whereby the SMC-MUNE procedure obtained two different model fits; MU1+MU4 in R10 and MU2+MU3 in R50. As a consequence, the estimated excitation range for MU1 in R50 is unusually large, leading to a relatively flat excitability curve with an enlarged coefficient of variation in relation to other MUs and between datasets. Nevertheless, the net effect of these firing combinations with the estimated expected MUTFs, see Figure 5 inserts, does not suggest that the overall description of the two

datasets greatly differ. This exemplifies the difficulty in differentiating between MUs with similar excitation and twitch characteristics. The difference in fit could have occurred in part due to the 70 mN response level in the R50 dataset not being represented within dataset R10.




Figure H.6: Median and 95% credible interval for the coefficient of variation for the random variable associated with the excitability curves presented in Figure 5, together with the mean and 95% confidence interval from Hales et al..

Table H.6: Most probable firing events (1=fire, 0=latent) for each level in the stimulus-response curve. The labelled MUs for R50 are re-ordered to demonstrate similarity between the two datasets. The response level around 70 mN is not present in the R10 dataset.

R10								R50										
Level (mN)	1	2	3	4	5	6	7	8	Level (mN)	1	2	3	4	5	6	7	8	
0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
50	1	0	0	0	0	0	0	0	40	1	0	0	0	0	0	0	0	0
–									70	0	1	0	0	0	0	0	0	0
100	1	1	0	0	0	0	0	0	110	1	1	0	0	0	0	0	0	0
120	1	0	0	1	0	0	0	0	120	0	1	0	1	0	0	0	0	0
170	1	1	0	1	0	0	0	0	150	1	1	0	1	0	0	0	0	0
210	1	1	1	1	0	0	0	0	200	1	1	1	1	0	0	0	0	0
230	1	1	1	1	1	0	0	0	230	1	1	1	1	1	0	0	0	0
270	1	1	1	1	1	1	0	0	270	1	1	1	1	1	1	0	0	0
320	1	1	1	1	1	1	1	0	300	1	1	1	1	1	1	1	0	0
360	1	1	1	1	1	1	1	1	350	1	1	1	1	1	1	1	1	1

References

- Aitkin, M., 2001. Likelihood and bayesian analysis of mixtures. *Statistical Modelling* 1, 287–304.
- Andrieu, C., 2007. Discussion of: Motor unit number estimation using reversible jump Markov chain Monte Carlo methods. *J. R. Stat. Soc. Ser. C* 56, 261–263.
- Bostock, H., 1983. The strength-duration relationship for excitation of myelinated nerve: computed dependence on membrane parameters. *Journal of Physiology* 341, 59–74.
- Bromberg, M.B., 2003. Consensus. *Supplements to Clinical Neurophysiology* 55, 333–338.
- Bromberg, M.B., 2007. Updating motor unit number estimation (MUNE). *Clinical Neurophysiology* 118, 1–8.
- Brown, W.F., Milner-Brown, H.S., 1976. Some electrical properties of motor units and their effects on the methods of estimating motor unit numbers. *Journal of Neurology, Neurosurgery & Psychiatry* 39, 249–257.
- Bungartz, H.J., Dirnstorfer, S., 2003. Multivariate quadrature on adaptive sparse grids. *Computing* 71, 89–114.
- Carvalho, C.M., Johannes, M.S., Lopes, H.F., Polson, N.G., 2010. Particle learning and smoothing. *Statistical Science. A Review Journal of the Institute of Mathematical Statistics* 25, 88–106.
- Casella, G.T.B., Almeida, V.W., Grumbles, R.M., Liu, Y., Thomas, C.K., 2010. Neurotrophic factors improve muscle reinnervation from embryonic neurons. *Muscle & Nerve* 42, 788–797.
- Daube, J.R., 1995. Estimating the number of motor units in a muscle. *Journal of Clinical Neurophysiology* 12, 585–594.
- Drovandi, C.C., Pettitt, A.N., Henderson, R.D., McCombe, P.A., 2014. Marginal reversible jump Markov chain Monte Carlo with application to motor unit number estimation. *Comput. Stat. Data Anal.* 72, 128–146.
- Genz, A., Bretz, F., 2009. Computation of multivariate normal and t probabilities. volume 195 of *Lecture Notes in Statistics*. Springer, Dordrecht.
- Gooch, C.L., Doherty, T.J., Chan, K.M., Bromberg, M.B., Lewis, R.A., Stashuk, D.W., Berger, M.J., Andary, M.T., Daube, J.R., 2014. Motor unit number estimation: a technology and literature review. *Muscle & Nerve* 50, 884–893.
- Green, P.J., 1995. Reversible jump Markov chain Monte Carlo computation and Bayesian model determination. *Biometrika* 82, 711–732.

- Hales, J.P., Lin, C.S.Y., Bostock, H., 2004. Variations in excitability of single human motor axons, related to stochastic properties of nodal sodium channels. *Journal of Physiology* 559, 953–964.
- Henderson, R.D., Ridall, G.R., Pettitt, A.N., McCombe, P.A., Daube, J.R., 2006. The stimulus–response curve and motor unit variability in normal subjects and subjects with amyotrophic lateral sclerosis. *Muscle & Nerve* 34, 34–43.
- Hol, J.D., Schon, T.B., Gustafsson, F., 2006. On resampling algorithms for particle filters, in: *Nonlinear Statistical Signal Processing Workshop, 2006 IEEE, IEEE*. pp. 79–82.
- Kadrie, H.A., Yates, S.K., Milner-Brown, H.S., Brown, W.F., 1976. Multiple point electrical stimulation of ulnar and median nerves. *Journal of Neurology, Neurosurgery & Psychiatry* 39, 973–985.
- Major, L.A., Hegedus, J., Weber, D.J., Gordon, T., Jones, K.E., 2007. Method for counting motor units in mice and validation using a mathematical model. *Journal of Neurophysiology* 97, 1846–1856.
- Major, L.A., Jones, K.E., 2005. Simulations of motor unit number estimation techniques. *Journal of Neural Engineering* 2, 17.
- McComas, A.J., Fawcett, P., Campbell, M.J., Sica, R.E.P., 1971. Electrophysiological estimation of the number of motor units within a human muscle. *Journal of Neurology, Neurosurgery & Psychiatry* 34, 121–131.
- Neal, R., 2003. Slice sampling. *The Annals of Statistics* 31, 705–767.
- Neal, R., 2011. MCMC using Hamiltonian dynamics, in: Brooks, S., Gelman, A., Jones, G., Meng, X.L. (Eds.), *The Handbook of Markov chain Monte Carlo*. Chapman and Hall/CRC, Boca Raton, London, New York. chapter 5, pp. 113–162.
- Pitt, M.K., Shephard, N., 1999. Filtering via simulation: auxiliary particle filters. *J. Amer. Statist. Assoc.* 94, 590–599.
- Ridall, P.G., Pettitt, A.N., Friel, N., McCombe, P.A., Henderson, R.D., 2007. Motor unit number estimation using reversible jump Markov chain Monte Carlo methods. *J. R. Stat. Soc. Ser. C* 56, 235–269.
- Ridall, P.G., Pettitt, A.N., Henderson, R.D., McCombe, P.A., 2006. Motor unit number estimation—a Bayesian approach. *Biometrics* 62, 1235–1250.
- Shefner, J.M., Cudkowicz, M., Brown, R.H., 2006. Motor unit number estimation predicts disease onset and survival in a transgenic mouse model of amyotrophic lateral sclerosis. *Muscle & Nerve* 34, 603–607.

Stashuk, D.W., Doherty, T.J., Kassam, A., Brown, W.F., 1994. Motor unit number estimates based on the automated analysis of F-responses. *Muscle & Nerve* 17, 881–890.

Stephens, M., 2000. Bayesian analysis of mixture models with an unknown number of components — an alternative to reversible jump methods. *The Annals of Statistics* 28, 40–74.