Using Monte Carlo Particle Methods to Estimate and Quantify Uncertainty in Periodic Parameters

Andrea Arnold

Abstract Estimating and quantifying uncertainty in system parameters remains a big challenge in applied and computational mathematics. A subset of these problems includes estimating periodic parameters that have unknown dynamics. Along with their time series, the period of these parameters may also be unknown and need to be estimated. The aim of this paper is to address the periodic parameter estimation problem, with particular focus on exploring the associated uncertainty, using Monte Carlo particle methods, such as the ensemble Kalman filter. Both parameter tracking and piecewise function approximations of periodic parameters are considered, highlighting aspects of parameter uncertainty in each approach when considering factors such as the frequency of available data and the number of piecewise segments used in the approximation. Estimation of the period of the periodic parameters and related uncertainty is also analyzed in the piecewise formulation. The pros and cons of each approach are discussed relative to a numerical example estimating the external voltage parameter in the FitzHugh-Nagumo system for modeling the spiking dynamics of neurons.

Key words: Particle methods; sequential Monte Carlo; ensemble Kalman filter; parameter estimation; uncertainty quantification; periodic parameters.

1 Introduction

Estimating and quantifying uncertainty in system parameters remains a big challenge in applied and computational mathematics. A subset of these problems includes estimating parameters that vary periodically with time but have unknown or uncertain time evolution models. Examples of periodic, time-varying parameters in dynamical

---

Andrea Arnold
Department of Mathematical Sciences, Worcester Polytechnic Institute, Worcester, MA, USA,
e-mail: anarnold@wpi.edu

systems arising from life sciences applications include the seasonal transmission in modeling the spread of infectious diseases [1, 9, 17] and the external voltage in modeling the spiking dynamics of neurons [34].

While most traditional algorithms aim at estimating constant parameters, the challenge in estimating time-varying parameters lies in accurately accounting for their time evolution without observations or known evolution models. In the case of periodic parameters, the resulting time series estimates should also maintain periodicity. Along with their time series, the period of these parameters may also be unknown and therefore may need to be estimated. This is particularly true in real data applications where a reasonable approximation of the period may not be clear from the available information.

The aim of this paper is to address the periodic parameter estimation problem, with particular focus on exploring the uncertainty associated with estimating periodic, time-varying parameters. In particular, this work uses sequential Monte Carlo particle methods (or nonlinear filtering methods) [27, 26, 16, 13, 29] to estimate the time series of periodic parameters. Note that while the term "sequential Monte Carlo" sometimes refers exclusively to particle filters, in this work the term more generally refers to sequential-in-time, Monte Carlo-based particle methods, including both particle filters and ensemble Kalman-type filters. In the Bayesian family of parameter estimation algorithms, Monte Carlo particle methods naturally account for uncertainty in the resulting parameter estimates by treating the unknowns as random variables with probability distributions describing their most likely values.

Both parameter tracking [34, 20, 30] and piecewise function approximations [8] of periodic parameters are considered, highlighting aspects of parameter uncertainty in each approach when considering factors such as the frequency of available data and the number of piecewise segments used in the approximation. Estimation of the period of the periodic parameters and related uncertainty is also analyzed in the piecewise formulation. As is demonstrated in the numerical results, while the parameter tracking method is efficient in tracking the overall behavior of slowly-varying parameters, it is unable to guarantee that periodicity is maintained in resulting parameter estimates. Pros and cons of each approach are discussed as applied to a numerical example estimating the external voltage parameter in the FitzHugh-Nagumo system for modeling neuron spiking dynamics.

The paper is organized as follows. Section 2 gives a review of the parameter estimation inverse problem and the Bayesian solution using sequential Monte Carlo particle methods, specifically outlining the augmented ensemble Kalman filter. Section 3 describes the parameter tracking and piecewise function approaches to estimating periodic parameters and discusses aspects of uncertainty relating to each approach. Section 4 gives numerical results on estimating the external voltage parameter in the FitzHugh-Nagumo model, and Section 5 provides discussion and future work.

2 Parameter Estimation and Monte Carlo Particle Methods

The parameter estimation inverse problem can be summarized as estimating unknown or uncertain system parameters given some discrete, noisy observations of (possibly a subset or some function of) the states of the system. More specifically, assume that an ordinary differential equation (ODE) model of the form

$$\frac{dx}{dt} = f(t, x, \theta), \qquad x(0) = x_0 \tag{1}$$

describes the dynamics of a system, which involves states $x = x(t) \in \mathbb{R}^d$ and unknown (or poorly known) parameters $\theta \in \mathbb{R}^q$. While the model function $f : \mathbb{R} \times \mathbb{R}^d \times \mathbb{R}^q \to \mathbb{R}^d$ is assumed to be known, the initial value $x_0 \in \mathbb{R}^d$ may also be unknown – in this case, x_0 may also be estimated along with the parameters θ. Further, assume the discrete, noisy observations $y_k \in \mathbb{R}^m$, $k = 1, 2, ..., T$, have the form

$$y_k = g(x(t_k), \theta) + w_k, \qquad 0 < t_1 < t_2 < \ldots < t_T \tag{2}$$

where $g : \mathbb{R}^d \times \mathbb{R}^q \to \mathbb{R}^m$, $m \leq d$, is a known observation function and w_k represents the observation error. The inverse problem is therefore to estimate the parameters θ and states $x(t)$ at some discrete times from the observations y_k.

From the Bayesian perspective, the unknown parameters θ, states x, and observations y are treated as random variables with probability distributions $\pi(\cdot)$, and the solution to the inverse problem is the joint posterior density

$$\pi(x, \theta \mid y) \propto \pi(y \mid x, \theta)\pi(x, \theta) \tag{3}$$

which follows from Bayes' theorem. The likelihood $\pi(y \mid x, \theta)$ indicates how likely it is that the data y are observed if the states x and parameters θ were known, and the prior density $\pi(x, \theta)$ encodes any information known about the states and parameters before accounting for the data.

There are various approaches to solving Bayesian inverse problems, including both sequential and nonsequential methods. Nonsequential methods, such as Markov chain Monte Carlo (MCMC)-type schemes [3, 19, 18], sample the posterior density by taking into account the full time series of data at once. Sequential Monte Carlo particle methods [27, 26, 16], on the other hand, make use of stochastic evolution-observation models to sequentially update the posterior using a two-step, predictor-corrector-type scheme, accounting for each data point as it arrives in time. A variety of Monte Carlo particle methods are available in the literature, including particle filters [29, 32, 24, 6] and ensemble Kalman-type filters [12, 11, 13, 7]. For a recent review, see [14].

Given the set $D_k = \{y_1, y_2, \ldots, y_k\}$ of observations up to time t_k, sequential Monte Carlo particle methods update the posterior distribution from time t_k to time t_{k+1} as follows:

$$\pi(x_k, \theta \mid D_k) \longrightarrow \pi(x_{k+1}, \theta \mid D_k) \longrightarrow \pi(x_{k+1}, \theta \mid D_{k+1}) \tag{4}$$

The first step (i.e., the prediction step) in the scheme predicts the values of the states at time t_{k+1} without knowledge of the data, while the second step (i.e., the analysis step) updates the predictions by taking into account the data at time t_{k+1}. Note that if there is no data observed at t_{k+1}, then $D_{k+1} = D_k$ and the prediction density $\pi(x_{k+1}, \theta \mid D_k)$ is equivalent to the posterior $\pi(x_{k+1}, \theta \mid D_{k+1})$. Starting with a prior density $\pi(x_0, \theta_0 \mid D_0)$, $D_0 = \emptyset$, this updating scheme is repeated until the final posterior density is obtained when $k = T$.

2.1 Augmented Ensemble Kalman Filter

The ensemble Kalman filter (EnKF) [12, 11] is a sequential particle approach that, unlike other particle methods that require importance sampling, moves (or pushes) particles forward in time based on the prediction and correction steps of the filter. Assume that the current density $\pi(x_k, \theta \mid D_k)$ is represented by a discrete ensemble

$$\mathcal{S}_{k|k} = \left\{ \left(x^n_{k|k}, \theta^n_{k|k} \right) \right\}_{n=1}^{N} \tag{5}$$

comprising N joint samples of the states $x^n_{k|k}$ and parameters $\theta^n_{k|k}$ at time k. In the prediction step of the filter, the state ensemble is updated using the equation

$$x^n_{k+1|k} = F(x^n_{k|k}, \theta^n_{k|k}) + v^n_{k+1}, \qquad v^n_{k+1} \sim \mathcal{N}(0, \mathsf{C}_{k+1}) \tag{6}$$

for each $n = 1, \ldots, N$, where F is the numerical solution to the ODE system (1) from time k to $k+1$. Note that the parameter samples $\theta^n_{k|k}$ are not updated in the prediction step.

To prepare for the analysis step, in which both the states and parameter values will be updated, the predicted state ensemble is combined with the current parameter ensemble into the augmented vectors

$$z^n_{k+1|k} = \begin{bmatrix} x^n_{k+1|k} \\ \theta^n_{k|k} \end{bmatrix} \in \mathbb{R}^{d+q}, \qquad n = 1, \ldots, N \tag{7}$$

and ensemble statistics formulas are used to compute the augmented ensemble mean $\bar{z}_{k+1|k}$ and covariance $\Gamma_{k+1|k}$. The covariance matrix $\Gamma_{k+1|k}$ contains cross-correlation information between the states and parameters that is used to update the parameter values in the next step.

In the analysis step, an observation ensemble

$$y^n_{k+1} = y_{k+1} + w^n_{k+1}, \qquad w^n_{k+1} \sim \mathcal{N}(0, \mathsf{D}_{k+1}), \qquad n = 1, \ldots, N \tag{8}$$

is generated around the observation y_{k+1} to prevent the resulting posterior ensemble from having too low a variance [11]. The observation ensemble is then compared to the observation model predictions

$$\widehat{y}^n_{k+1} = g(x^n_{k+1|k}, \theta^n_{k|k}), \qquad n = 1, \ldots, N \tag{9}$$

with g as in (2) in the updating equation

$$z^n_{k+1|k+1} = z^n_{k+1|k} + \mathsf{K}_{k+1}\left(y^n_{k+1} - \widehat{y}^n_{k+1}\right), \qquad n = 1, \ldots, N. \tag{10}$$

To accommodate nonlinear observations [31], the Kalman gain K_{k+1} in (10) is computed by

$$\mathsf{K}_{k+1} = \mathsf{S}^{z\widehat{y}}_{k+1}\left(\mathsf{S}^{\widehat{y}\widehat{y}}_{k+1} + \mathsf{D}_{k+1}\right)^{-1} \tag{11}$$

where $\mathsf{S}^{z\widehat{y}}_{k+1}$ gives the cross-correlation between the augmented predictions $z^n_{k+1|k}$ in (7) and observation model predictions $\widehat{y}^n_{k+1}$ in (9), $\mathsf{S}^{\widehat{y}\widehat{y}}_{k+1}$ is the forecast error covariance, and D_{k+1} is the observation noise covariance as in (8). The above algorithm, known as the augmented EnKF for combined state and parameter estimation [13, 7], is repeated until the joint posterior density is obtained at $k = T$.

3 Estimating Periodic Parameters and the Role of Uncertainty

In the traditional Monte Carlo particle methods described in Section 2, the parameters θ are assumed to be constant (or static) parameters, i.e., $d\theta/dt = 0$, and are artificially evolved over time as the posterior is updated. Depending on the implementation of the method used, the parameter values may be updated during both the prediction and analysis steps, or only in the analysis step via their correlation with the state predictions. In particular, the augmented EnKF outlined in Section 2.1 updates the parameter estimates only in the analysis step at each data arrival through the use of cross-correlation information encoded in the Kalman gain (11). The periodic parameters of interest in this work, however, are known to vary with time but do not have known time evolution models. The main challenges in this problem therefore lie in accurately accounting for the time evolution of these parameters while also maintaining the periodic structure.

One approach is to consider parameter tracking algorithms [34, 20, 30], which can trace the dynamics of slowly-changing parameters over time by allowing for a drift in the parameter values during the prediction step of sequential Monte Carlo. More specifically, the predicted change in the parameter $\theta(t)$ is modeled as a random walk

$$\theta^n_{k+1|k} = \theta^n_{k|k} + \xi^n_{k+1}, \qquad \xi^n_{k+1} \sim \mathcal{N}(0, \mathsf{E}_{k+1}), \qquad n = 1, \ldots, N \tag{12}$$

where E_{k+1} defines the covariance of the drift term ξ^n_{k+1}. Note that inclusion of the drift term in (12) is crucial in allowing the algorithm to track the underlying dynamics of the time-varying parameter. While parameter tracking algorithms are straightforward to implement, the drift covariance E_{k+1}, which is typically modeled as $\mathsf{E}_{k+1} = \sigma^2_\xi \mathsf{I}$ for some constant σ_ξ, must be chosen carefully in order to avoid filter divergence [23, 2, 37, 10, 21] and result in a useful parameter estimate. The drift covariance also plays a direct role in the uncertainty of the resulting parameter estimate, thereby affecting the corresponding model output predictions [5]. Moreover, in the case of estimating periodic parameters, parameter tracking algorithms do not guarantee that periodicity is maintained throughout the estimation process.

An alternative approach that maintains periodicity in the parameter estimation is to model the periodic parameter $\theta(t)$ as a piecewise function

$$\theta(t) = \begin{cases} \theta_1(t)\,, & t \in \left[0, \frac{p}{\ell}\right) \\ \theta_2(t)\,, & t \in \left[\frac{p}{\ell}, \frac{2p}{\ell}\right) \\ \vdots & \vdots \\ \theta_\ell(t)\,, & t \in \left[\frac{(\ell-1)p}{\ell}, p\right) \end{cases} \tag{13}$$

where each $\theta_i(t)$, $i = 1, \ldots, \ell$, is a function relying on some unknown constant coefficients, repeated each period p, that can be estimated using traditional Monte Carlo particle methods. A similar piecewise formulation using nonlinear filtering was presented in [8] assuming that the period p was known and fixed during the estimation process. However, in general the period of the parameter may not be known a priori and may need to be estimated along with the other unknown system parameters. Therefore, in this work, the period p is assumed to be unknown and is estimated along with the unknown piecewise function coefficients.

The formulation in (13) can accommodate estimation using piecewise constant functions or splines of various order. In this study and the numerical experiments that follow in Section 4, we employ a continuous linear interpolating spline (of degree 1) where

$$\theta_i(t) = a_i + b_i(t - t_{i-1})\,, \quad t \in [t_{i-1}, t_i) = \left[\frac{(i-1)p}{\ell}, \frac{ip}{\ell}\right) \tag{14}$$

for $i = 1, \ldots, \ell$, with constant coefficients a_i and b_i denoting the y-intercept and slope of the line $\theta_i(t)$, respectively. Note that the spline knots t_j, $j = 0, \ldots, \ell$, in (14) depend on both the period p and number of spline segments ℓ. Continuity dictates that $\theta_i(t_i) = \theta_{i+1}(t_i)$ for $i = 1, \ldots, \ell - 1$, and it follows from definition of the linear spline in (14) that $\theta_i(t_{i-1}) = a_i$ for $i = 1, \ldots, \ell$. Since the slope coefficients b_i can be computed directly from the y-intercepts a_i, $i = 1, \ldots, \ell + 1$, via the formula

$$b_i = \frac{a_{i+1} - a_i}{t_i - t_{i-1}} = \frac{a_{i+1} - a_i}{p/\ell} = \frac{\ell}{p}\left(a_{i+1} - a_i\right) \tag{15}$$

it suffices to estimate only the values for a_i, $i = 1, \ldots, \ell + 1$, along with the period p. Therefore the parameter estimation problem consists of estimating $L = \ell + 1$ spline coefficients and the period p, for a total of $L + 1$ unknown static parameters relating to the periodic parameter of interest.

Various factors must be considered in analyzing the uncertainty relating to the piecewise formulation (13)–(14). In this study, we consider how the frequency of the data in time and the number of linear spline segments ℓ used in the estimation affects the resulting periodic parameter estimates and corresponding uncertainty.

4 Numerical Example: External Voltage in FitzHugh-Nagumo

As a numerical example, we consider synthetic data generated from the FitzHugh-Nagumo system [15] which acts as a simplified version of the Hodgkin-Huxley system [22] for modeling the spiking dynamics of single neurons. The FitzHugh-Nagumo equations are given by

$$\frac{dx_1}{dt} = c\Big(x_2 + x_1 - \frac{x_1^3}{3} + v(t)\Big) \tag{16}$$

$$\frac{dx_2}{dt} = -\frac{1}{c}\Big(x_1 - a + bx_2\Big) \tag{17}$$

where the state variable $x_1(t)$ represents the measurable membrane potential of the neuron, while $x_2(t)$ denotes an unobservable combined effect of various ionic currents. The parameters a, b, and c are commonly fixed to some known values a priori, but the external voltage $v(t)$ is an unknown, time-varying parameter.

Figure 1 shows the synthetic data and underlying system states generated from (16)–(17) using initial values $x_1(0) = 1$ and $x_2(0) = 0.5$ and fixed parameters $a = 0.7$, $b = 0.8$, and $c = 3$, along with the time-varying external voltage parameter modeled as a periodic, sinusoidal function $v(t) = 0.5 \sin(\omega t + \pi/2) - 1$ with frequency $\omega = 0.1$. Therefore, in this example, $v(t)$ plays the role of the periodic parameter $\theta(t)$ described in Section 3. This choice of $v(t)$ varies more slowly than the system dynamics, making it amenable to particle methods with parameter tracking. A similar example was considered in [5], where the focus was to study the effects of uncertainty in parameter tracking estimates and their corresponding model output predictions. The data was generated by observing $x_1(t)$ at 1257 equidistant time instances over the interval $[0, 251.2]$, covering four periods of $v(t)$, and corrupting the observations with zero-mean Gaussian noise. The standard deviation of the noise was taken to be 20% of the standard deviation of $x_1(t)$ over the full time interval.

For the first numerical experiment, we consider estimating the periodic parameter $v(t)$ using the piecewise formulation (13)–(14) with $\ell = 10$ spline segments ($L = 11$ knots) and estimating the $L = 11$ unknown linear spline coefficients $a_1, \dots, a_{11}$, along with the unknown period p. While various particle methods could be applied, we employ an augmented EnKF in the style of [7] with $N = 150$ ensemble members to estimate the system states $x_1(t)$ and $x_2(t)$ along with the parameter vector $\theta = (a_1, \dots, a_{11}, p) \in \mathbb{R}^{12}$ as described; see [7] for implementation details of the filter beyond those given in Section 2.1.

Assuming that the initial values of the system are not fully known, the prior ensemble of states is drawn uniformly from 0.5 to 1.5 times the value of the first observed value of $x_1(t)$ and set to 0 for $x_2(t)$ (unobserved). The prior ensemble of

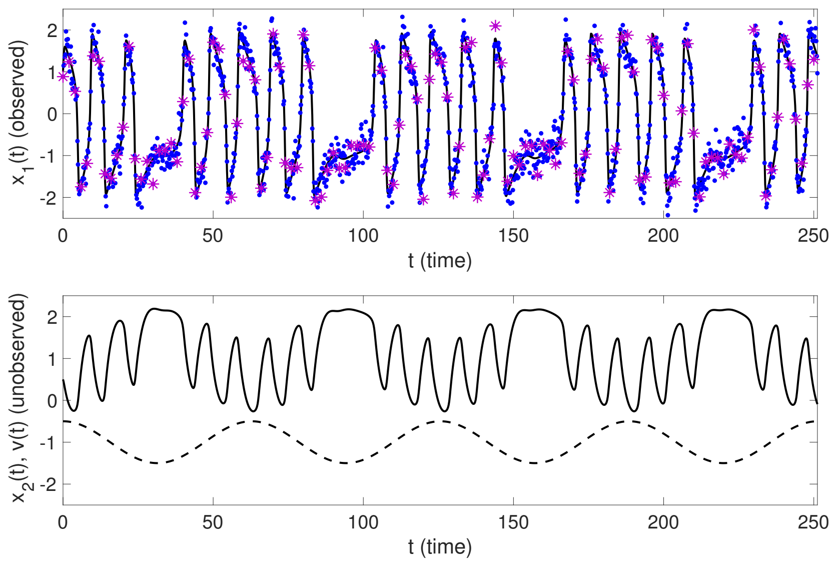

Fig. 1 Noisy observations of the membrane potential $x_1(t)$ (top, blue and purple markers) from the FitzHugh-Nagumo system (16)–(17), along with the unobserved lumped ionic current $x_2(t)$ (bottom, solid black) and external voltage parameter $v(t)$ (bottom, dashed black). In the top panel, the blue (dots) and purple (asterisks) markers together represent noisy observations taken every 0.2 time units, while the purple markers alone show noisy observations every 2 time units.

parameter values is drawn uniformly from $\mathcal{U}(-2, 1)$ for each of the spline coefficients $a_1, \ldots, a_{11}$ and from $\mathcal{U}(55, 75)$ for the period p. Throughout the estimation process, a positivity constraint is placed on the period such that $p^n_{k|k} > 0$ for all n and k, and time integration is performed using the Adams-Moulton linear multistep methods of orders 1 and 2 [28, 25].

Figure 2 shows the resulting linear spline estimates of $v(t)$ computed using the estimated parameter means and ± 2 standard deviation values for the spline coefficients $a_1, \ldots, a_{11}$ and period p, repeated over four periods. The corresponding estimates of p are listed in Table 1. For comparison with the piecewise approach, Figure 2 also shows the resulting mean and ± 2 standard deviation curves estimating $v(t)$ using the augmented EnKF with parameter tracking, where the drift term has prescribed standard deviation $\sigma_\xi = 0.01$. Note that the uncertainty in the resulting estimates of $v(t)$ is much smaller in the piecewise formulation; however, some parts of the true $v(t)$ curve are not captured within the uncertainty bounds, specifically near the beginning of the estimated period (≈ 63.7212) at each repetition. While the uncertainty in the parameter drift estimate is able to almost fully capture the underlying true $v(t)$, the mean estimate does not fully maintain the periodicity intrinsic to $v(t)$.

The next numerical experiment explores how the uncertainty in the parameter estimates using both the piecewise linear spline formulation and parameter tracking is affected by the frequency of available data over time. To that end, the synthetic

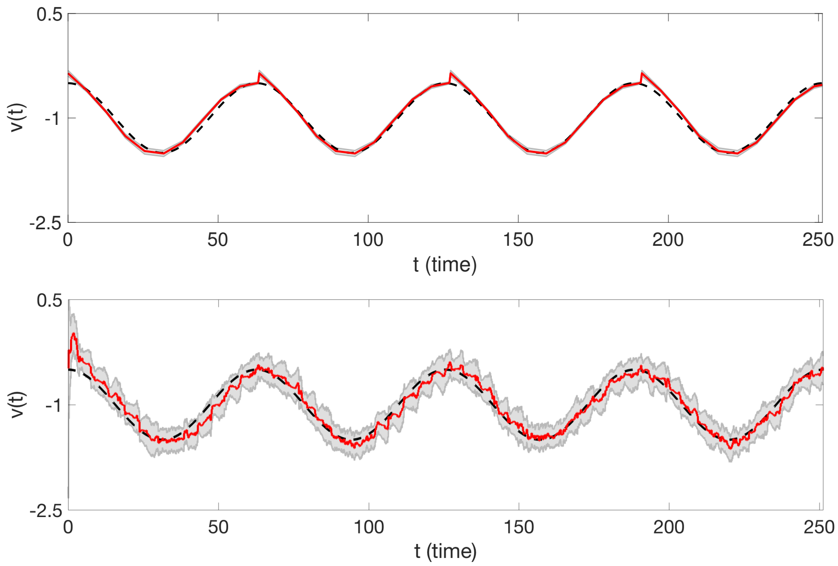

Fig. 2 Parameter estimates of the external voltage parameter $v(t)$ in the FitzHugh-Nagumo system (16)–(17) computed using piecewise linear splines with estimated spline coefficients when $\ell = 10$ and period p repeated over four periods (top panel) and parameter tracking (bottom panel). In the top panel, the linear spline using the augmented EnKF mean estimates of the spline coefficients and period is shown in solid red, while the linear splines computing using the ±2 standard deviation parameter estimates are shown in dark grey, filled with light grey. In the bottom panel, the mean parameter tracking estimating using the augmented EnKF is shown in solid red, while the ±2 standard deviation curves are shown in dark grey, filled with light grey. In both panels, the true $v(t)$ used in generating the synthetic data is shown in dashed black. Estimates were obtained using the full synthetic data shown in Figure 1.

data is subsampled, taking every 10 data points for a total of 126 noisy observations of $x_1(t)$ at equidistant time instances over the interval $[0, 251.2]$. Figure 1 displays the subsampled data in purple markers (asterisks) on the top panel. Figure 3 shows the resulting parameter estimates, using both piecewise linear splines and parameter tracking, initialized as in the previous numerical experiment. Note that less frequent observations result in significantly more uncertainty in both the piecewise linear spline and parameter tracking estimates of $v(t)$. While the linear spline estimate is able to fairly well approximate and fully capture the true $v(t)$ within the uncertainty bounds, the parameter tracking algorithm has more difficulty tracking $v(t)$ in this case – the mean estimate does not maintain periodicity and is also noticeably out of phase with the true $v(t)$.

The last numerical experiment considered in this paper studies the effect of the number of linear spline segments ℓ (corresponding to $L = \ell + 1$ spline knots) on the piecewise estimation of $v(t)$ and corresponding estimate for the period p. To this end, the piecewise linear spline estimation is performed using five different choices of ℓ (namely, $\ell = 2, 5, 10, 15, 20$) and the full synthetic data in Figure 1. Figure 4 shows

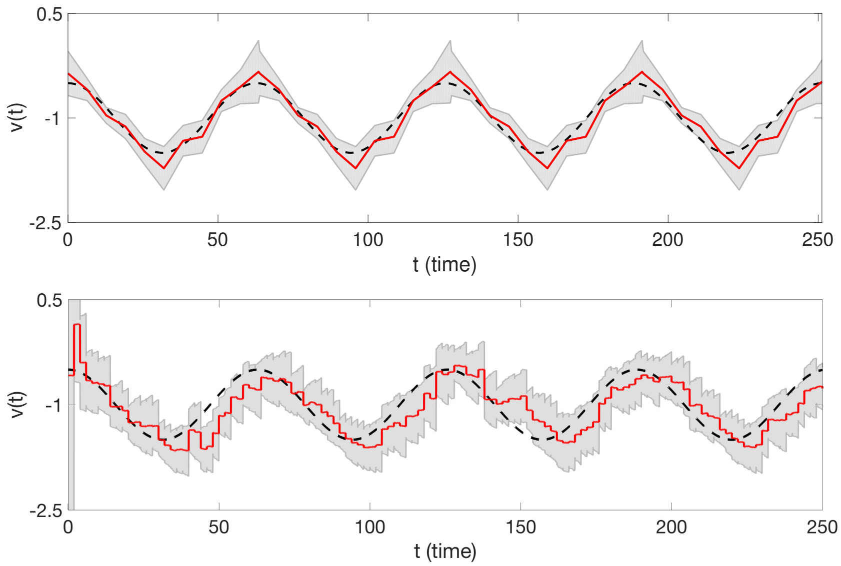

Fig. 3 Parameter estimates of the external voltage parameter $v(t)$ in the FitzHugh-Nagumo system (16)–(17) computed using piecewise linear splines with estimated spline coefficients when $\ell = 10$ and period p repeated over four periods (top panel) and parameter tracking (bottom panel). In the top panel, the linear spline using the augmented EnKF mean estimates of the spline coefficients and period is shown in solid red, while the linear splines computing using the ± 2 standard deviation parameter estimates are shown in dark grey, filled with light grey. In the bottom panel, the mean parameter tracking estimating using the augmented EnKF is shown in solid red, while the ± 2 standard deviation curves are shown in dark grey, filled with light grey. In both panels, the true $v(t)$ used in generating the synthetic data is shown in dashed black. Estimates were obtained using the subsampled synthetic data shown in Figure 1.

the resulting linear spline estimates for each ℓ over one estimated period. Table 1 gives the corresponding estimates of the period p in each case, along with the relative error comparing the EnKF mean estimate of the period with the true period used in generating the synthetic data. The relative error in each case is computed via the formula

$$\text{relative error} = \left| \frac{p_{\text{true}} - p_{\text{est}}}{p_{\text{true}}} \right| \tag{18}$$

where p_{est} is the augmented EnKF mean estimate of the period and $p_{\text{true}} \approx 62.8319$ is the true period (up to four decimal places).

The amount of uncertainty in each spline estimate is low, similar to the results seen in Figure 2 using the full time series of data. It is interesting to note that as the number of spline segments ℓ increases, the EnKF estimate of the period p tends to increase. While the fit of the spline improves from $\ell = 2$ to $\ell = 10$, adding more spline segments eventually starts to degrade the fit along with overestimating the period, as is the case when $\ell = 15$ and $\ell = 20$.

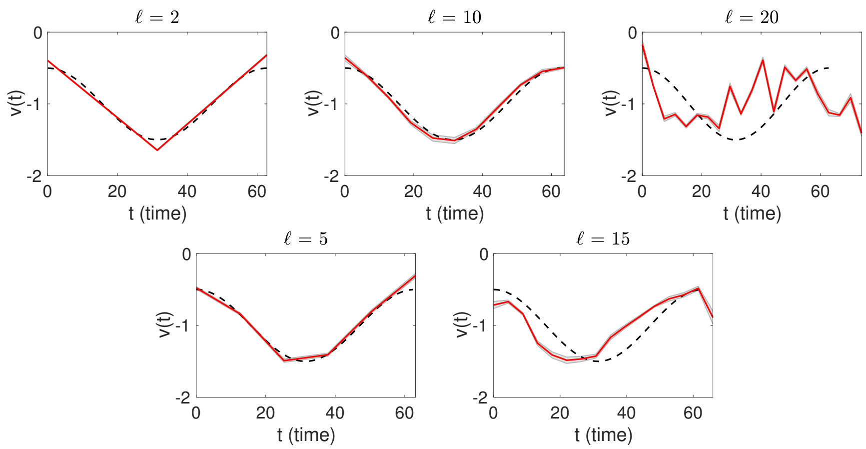

Fig. 4 Parameter estimates of the external voltage parameter $v(t)$ in the FitzHugh-Nagumo system (16)–(17) computed using piecewise linear splines with estimated spline coefficients when ℓ = 2, 5, 10, 15 and 20, respectively, and period p, shown over one period. In each panel, the linear spline using the augmented EnKF mean estimates of the spline coefficients and period is shown in solid red, while the linear splines computing using the ±2 standard deviation parameter estimates are shown in dark grey, filled with light grey. The true $v(t)$ used in generating the synthetic data is shown in dashed black. Estimates were obtained using the full synthetic data shown in Figure 1. Corresponding period estimates are given in Table 1.

Table 1 Augmented EnKF mean and ±2 standard deviation parameter estimates of the period p of the piecewise linear spline estimate of the external voltage parameter $v(t)$ for different numbers of spline segments ℓ. The relative error between the mean estimate and true value of p in each case is computed using the formula in (18). Values in the table are reported to four decimal places.

# of Spline Segments	Estimated p (mean ± 2 std)	Relative error (mean)
$\ell = 2$	62.7699 ± 0.0523	0.0009
$\ell = 5$	63.1169 ± 0.0620	0.0045
$\ell = 10$	63.7212 ± 0.0881	0.0142
$\ell = 15$	65.9231 ± 0.0540	0.0492
$\ell = 20$	73.8291 ± 0.0189	0.1750

5 Discussion

This paper addresses the problem of estimating and quantifying uncertainty in periodic, time-varying parameters using sequential Monte Carlo parameter estimation techniques. Estimation approaches using both particle methods with parameter tracking and piecewise linear spline (with spline coefficients and periods estimated using particle methods) are considered, and the role of uncertainty is highlighted in each. In particular, uncertainty relating to the frequency of available time series data and the number of spline segments used in the linear spline estimates is tested via numerical experiment on an electrophysiology example estimating the external voltage parameter in the FitzHugh-Nagumo system for modeling the spiking dynamics of single neurons.

As demonstrated in the numerical results in Section 4, there are pros and cons to using each of the presented approaches for estimating periodic parameters. One clear computational advantage of the parameter tracking algorithm is its straightforward implementation in the sequential Monte Carlo framework and flexibility in approximating the shape of the parameter of interest. Nothing is assumed about periodicity a priori, and there is only one parameter to track over time. However, since nothing is assumed about periodicity in the parameter tracking, the periodicity of the parameter is therefore neglected in the estimation process and periodicity is not maintained. The choice of the drift variance also has a significant impact on the resulting parameter tracking estimate, in terms of both accuracy and uncertainty. Moreover, the numerical experiments show that the parameter tracking algorithm has more difficulty tracking the periodic parameter as less frequent time series data is available.

The piecewise linear spline formulation maintains periodicity by prescribing a periodic form to the parameter a priori, then estimating the coefficients and period that best fit the available data via a particle approach. The numerical experiments show that the frequency of available time series data has a direct impact on the uncertainty relating to the linear spline estimates, with more frequent data resulting in tighter uncertainty bounds. The number of spline segments ℓ also has a significant effect on both the fit of the resulting spline and the corresponding period estimation. An interesting problem would be to consider estimating ℓ along with the period p; however, this is not straightforward, as ℓ and p depend on one another in the piecewise formulation in (13). Instead, one could interpret the problem of choosing ℓ as a model selection problem and could apply available methods for model selection; see, e.g., [36, 35, 4, 33]. This remains as future work.

Note that while the linear splines shown in Figures 2 and 3 are formulated to be continuous within a given period, the piecewise formulation in (13)–(14) does not necessarily guarantee continuity of the spline between periods. In order to maintain continuity between periods, an additional constraint that $a_1 = a_\ell$ is required (but not considered in this work). Regarding the period estimation, note that special care must be taken in the implementation to avoid, e.g., negative or inappropriate values being assigned for the period of the periodic parameter. A simple approach used in the numerical results in this paper is to apply a positivity constraint within the Monte Carlo particle algorithm to retain $p^n_{k|k} > 0$ for all n and k. Further, since periodic parameters are subset of all possible parameters in (1), additional constant parameters, such as the initial states of the dynamical system, may be estimated simultaneously in both the parameter tracking and piecewise formulations.

While it is possible to use a variety of parameter estimation techniques, the use of Monte Carlo particle methods in this work provides a natural framework for analyzing the time series data typically available in applications where time-varying parameters are relevant. Moreover, Monte Carlo methods provide a natural measure of uncertainty in the parameter estimation, which can be used for model prediction and uncertainty quantification. Future work includes the design and analysis of parameter tracking-type Monte Carlo particle algorithms that incorporate structural characteristics like periodicity into the sequential estimation without relying on a piecewise functional form for the time-varying parameters.

Acknowledgements This work is supported by the National Science Foundation under grant number NSF/DMS-1819203.

References

1. Altizer, S., Dobson, A., Hosseini, P., Hudson, P., Pascual, M., Rohani, P.: Seasonality and the dynamics of infectious diseases. Ecology Letters **9**, 467–484 (2006)
2. Anderson, J.L.: An ensemble adjustment Kalman filter for data assimilation. Mon Weather Rev **129**, 2884–2903 (2001)
3. Andrieu, C., Thoms, J.: A tutorial on adaptive MCMC. Statistics and Computing **18**(4), 343–373 (2008)
4. Arlot, S., Celisse, A.: A survey of cross-validation procedures for model selection. Statistics Surveys **4**, 40–79 (2010)
5. Arnold, A.: Exploring the effects of uncertainty in parameter tracking estimates for the time-varying external voltage parameter in the FitzHugh-Nagumo model. In: P. Nithiarasu, M. Ohta, M. Oshima (eds.) 6th International Conference on Computational and Mathematical Biomedical Engineering – CMBE2019, pp. 512–515 (2019)
6. Arnold, A., Calvetti, D., Somersalo, E.: Linear multistep methods, particle filtering and sequential Monte Carlo. Inverse Problems **29**(8), 085007 (2013)
7. Arnold, A., Calvetti, D., Somersalo, E.: Parameter estimation for stiff deterministic dynamical systems via ensemble Kalman filter. Inverse Problems **30**(10), 105008 (2014)
8. Arnold, A., Lloyd, A.L.: An approach to periodic, time-varying parameter estimation using nonlinear filtering. Inverse Problems **34**(10), 105005 (2018)
9. Aron, J., Schwartz, I.: Seasonality and period-doubling bifurcations in an epidemic model. J Theor Biol **110**, 665–679 (1984)
10. Berry, T., Sauer, T.: Adaptive ensemble Kalman filtering of non-linear systems. Tellus A **65**, 20331 (2013)
11. Burgers, G., van Leeuwen, P., Evensen, G.: Analysis scheme in the ensemble Kalman filter. Mon Weather Rev **126**(6), 1719–1724 (1998)
12. Evensen, G.: Sequential data assimilation with a nonlinear quasi-geostrophic model using Monte Carlo methods to forecast error statistics. J Geophys Res **99**(C5), 10143–10162 (1994)
13. Evensen, G.: The ensemble Kalman filter for combined state and parameter estimation. IEEE Control Syst Mag **29**(3), 83–104 (2009)
14. Fearnhead, P., Kunsch, H.R.: Particle filters and data assimilation. Annual Review of Statistics and Its Application **5**, 421–449 (2018)
15. FitzHugh, R.: Impulses and physiological states in theoretical models of nerve membrane. Biophys J **1**, 445–466 (1961)
16. Gordon, N.J., Salmond, D.J., Smith, A.F.M.: Novel approach to nonlinear/non-Gaussian Bayesian state estimation. IEEE Proceedings-F **140**(2), 107–113 (1993)
17. Grassly, N., Fraser, C.: Seasonal infectious disease epidemiology. Proc R Soc B **273**, 2541–2550 (2006)
18. Haario, H., Laine, M., Mira, A., Saksman, E.: DRAM: Efficient adaptive MCMC. Statistics and Computing **16**, 339–354 (2006)
19. Haario, H., Saksman, E., Tamminen, J.: An adaptive Metropolis algorithm. Bernoulli **7**, 223–242 (2001)
20. Hamilton, F., Berry, T., Peixoto, N., Sauer, T.: Real-time tracking of neuronal network structure using data assimilation. Physical Review E **88**, 052715 (2013)

21. Harlim, J., Majda, A.J.: Catastrophic filter divergence in filtering nonlinear dissipative systems. Commun Math Sci **8**(27-43) (2010)
22. Hodgkin, A., Huxley, A.: A quantitative description of membrane current and its application to conduction and excitation in nerve. J Physiol **117**, 500–544 (1952)
23. Houtekamer, P.L., Mitchell, H.L.: Data assimilation using an ensemble Kalman filter technique. Mon Weather Rev **126**, 796–811 (1998)
24. Ionides, E., Breto, C., King, A.: Inference for nonlinear dynamical systems. PNAS **103**(49), 18438–18443 (2006)
25. Iserles, A.: A First Course in the Numerical Analysis of Differential Equations, 2 edn. Cambridge Texts in Applied Mathematics. Cambridge University Press, New York (2009)
26. Kantas, N., Doucet, A., Singh, S.S., Maciejowski, J., Chopin, N.: On particle methods for parameter estimation in state-space models. Statistical Science **30**(3), 328–351 (2015)
27. Kitagawa, G.: A self-organizing state-space model. Journal of the American Statistical Association **93**(443), 1203–1215 (1998)
28. LeVeque, R.J.: Finite Difference Methods for Ordinary and Partial Differential Equations. SIAM, Philadelphia (2007)
29. Liu, J., West, M.: Combined parameter and state estimation in simulation-based filtering. In: A. Doucet, N. de Freitas, N. Gordon (eds.) Sequential Monte Carlo Methods in Practice, pp. 197–223. Springer, New York (2001)
30. Matzuka, B.: Nonlinear filtering methodologies for parameter estimation and uncertainty quantification in noisy, complex biological systems. Ph.D. thesis, North Carolina State University (2014)
31. Moradkhani, H., Sorooshian, S., Gupta, H., Houser, P.: Dual state-parameter estimation of hydrological models using ensemble Kalman filter. Adv Water Resour **28**(2), 135–147 (2005)
32. Pitt, M., Shephard, N.: Filtering via simulation: auxiliary particle filters. J Amer Statist Assoc **94**, 590–599 (1999)
33. Toni, T., Welch, D., Strelkowa, N., Ipsen, A., Stumpf, M.P.H.: Approximate Bayesian computation scheme for parameter inference and model selection in dynamical systems. J R Soc Interface **6**, 187–202 (2009)
34. Voss, H., Timmer, J., Kurths, J.: Nonlinear dynamical system identification from uncertain and indirect measurements. International Journal of Bifurcation and Chaos **14**(6), 1905–1933 (2004)
35. Vyshemirsky, V., Girolami, M.A.: Bayesian ranking of biochemical system models. Bioinformatics **24**(6), 833–839 (2007)
36. Wasserman, L.: Bayesian model selection and model averaging. Journal of Mathematical Psychology **44**, 92–107 (2000)
37. Whitaker, J.S., Hamill, T.M.: Ensemble data assimilation without perturbed observations. Mon Weather Rev **130**, 1913–1924 (2002)