

PAPER

An approach to periodic, time-varying parameter estimation using nonlinear filtering

To cite this article: Andrea Arnold and Alun L Lloyd 2018 *Inverse Problems* **34** 105005

Manuscript version: Accepted Manuscript

Accepted Manuscript is “the version of the article accepted for publication including all changes made as a result of the peer review process, and which may also include the addition to the article by IOP Publishing of a header, an article ID, a cover sheet and/or an ‘Accepted Manuscript’ watermark, but excluding any other editing, typesetting or other changes made by IOP Publishing and/or its licensors”

This Accepted Manuscript is© .

During the embargo period (the 12 month period from the publication of the Version of Record of this article), the Accepted Manuscript is fully protected by copyright and cannot be reused or reposted elsewhere.

As the Version of Record of this article is going to be / has been published on a subscription basis, this Accepted Manuscript will be available for reuse under a CC BY-NC-ND 3.0 licence after the 12 month embargo period.

After the embargo period, everyone is permitted to use copy and redistribute this article for non-commercial purposes only, provided that they adhere to all the terms of the licence <https://creativecommons.org/licenses/by-nc-nd/3.0>

Although reasonable endeavours have been taken to obtain all necessary permissions from third parties to include their copyrighted content within this article, their full citation and copyright line may not be present in this Accepted Manuscript version. Before using any content from this article, please refer to the Version of Record on IOPscience once published for full citation and copyright details, as permissions may be required. All third party content is fully copyright protected, unless specifically stated otherwise in the figure caption in the Version of Record.

View the [article online](#) for updates and enhancements.

An approach to periodic, time-varying parameter estimation using nonlinear filtering

Andrea Arnold^{1*,2,3} and Alun L. Lloyd^{2,3,4}

1 Department of Mathematical Sciences, Worcester Polytechnic Institute, Worcester, MA, USA
2 Department of Mathematics, North Carolina State University, Raleigh, NC, USA
3 Center for Quantitative Sciences in Biomedicine, North Carolina State University, Raleigh, NC, USA
4 Biomathematics Graduate Program, North Carolina State University, Raleigh, NC, USA

* corresponding author: anarnold@wpi.edu

Abstract

Many systems arising in biological applications are subject to periodic forcing. In these systems the forcing parameter is not only time-varying but also known to have a periodic structure. We present an approach to estimating periodic, time-varying parameters that imposes periodic structure by treating the time-varying parameter as a piecewise function with unknown coefficients. This method allows the resulting parameter estimate more flexibility in shape than prescribing a specific functional form (e.g., sinusoidal) to model its behavior, while still maintaining periodicity. We employ nonlinear filtering, more specifically, a version of the augmented ensemble Kalman filter (EnKF), to estimate the unknown coefficients comprising the piecewise approximation of the periodic, time-varying parameter. This allows for straightforward comparison of the proposed method with an EnKF-based parameter tracking algorithm, where periodicity is not guaranteed.

We demonstrate the effectiveness of the proposed approach on two biological examples: a synthetic example with data generated from the nonlinear FitzHugh-Nagumo system, modeling the excitability of a nerve cell, to estimate the external voltage parameter, and a case study using reported measles incidence data from three locations during the pre-vaccine era to estimate the seasonal transmission parameter. The formulation of the proposed approach also allows for simultaneous estimation of initial conditions and other static system parameters, such as the reporting probability of measles cases, which is vital for predicting under-reported incidence data.

Keywords: time-varying parameter estimation; periodic structure; nonlinear filtering; ensemble Kalman filter (EnKF); FitzHugh-Nagumo; measles transmission.

1 Introduction

Many systems arising in biological applications are subject to periodic forcing, such as the seasonal forcing seen in epidemiological systems [1–3] and the daily forcing in circadian rhythms [4–6]. In these types of systems, the forcing parameter is not only time-varying in nature but is also known to have a periodic structure. While it is possible to use periodic functions (e.g., sinusoids) to approximate the temporal behavior of such parameters [7, 8], the structural restrictions often do not adequately capture the true time evolution of the parameters. This illustrates the need for methodology to estimate time-varying parameters that is able to maintain known structural characteristics without imposing restrictive evolution models.

In this work we present an approach to estimating periodic, time-varying parameters using nonlinear filtering. In particular, we impose periodicity by treating the time-varying parameter as a piecewise function with unknown coefficients, repeated each period over the course of the time series, and then estimate the coefficients using a version of the ensemble Kalman filter (EnKF). Formulating the problem in this way allows the resulting time-varying parameter estimate to maintain its periodic structure without imposing any further restrictions to shape. The proposed approach also permits simultaneous estimation of unknown time-invariant parameters associated with the system, including initial conditions, as may be needed for certain applications.

Most parameter estimation methodology available in the literature is traditionally aimed at estimating time-invariant (static) parameters. Classical deterministic techniques for solving the static parameter estimation problem typically rely on least squares optimization routines [9–12]. Bayesian approaches include Markov chain Monte Carlo (MCMC) methods [13–16] and nonlinear filtering (or sequential Monte Carlo) methods such as particle filters [17–21] and ensemble Kalman-type filters [22–25].

In nonlinear filtering algorithms, static parameters are artificially evolved over time with the aim of converging to a constant value. If the true parameter values do change slowly over time (i.e., change on a scale slower than the dynamics of the system), then allowing the parameters to evolve with a random walk may capture the drift [26–28]. Parameter tracking with nonlinear filtering can be used to estimate periodic parameters, as in [26]. However, while being the least structurally restrictive method, parameter tracking via random walk evolution does not guarantee that the resulting time series estimate will maintain periodicity or any known structural characteristics inherent to the parameter. Further, special care must be taken in choosing the variance of the random walk in parameter tracking algorithms in order to avoid filter divergence [29–37] and obtain a useful parameter estimate.

Unlike in parameter tracking algorithms, where periodicity is not imposed, the formulation proposed in this work directly imposes periodic structure throughout the estimation process without forcing the parameter to have a sinusoidal shape. While the coefficients relating to the piecewise formulation could be estimated using various techniques, nonlinear filtering provides a natural framework to accommodate the sequential nature of the data and time-varying behavior of the parameters we aim to estimate. Moreover, use of nonlinear filtering to estimate the coefficients allows for straightforward comparison of the proposed approach with the parameter tracking algorithms as described.

The paper is organized as follows. In Section 2, we review the standard static parameter estimation inverse problem and its solution in the Bayesian statistical framework via nonlinear filtering methods, with particular focus on the augmented EnKF. After describing our proposed method for estimating periodic, time-varying parameters in Section 3, we demonstrate its effectiveness with two numerical examples in Section 4. We compare the proposed method to an EnKF-based parameter tracking algorithm, where periodicity is not guaranteed, using synthetic data generated

from the FitzHugh-Nagumo system, which describes the spiking dynamics of neurons. We further demonstrate the proposed method by estimating the seasonal transmission parameter in an epidemic model for the spread of measles. Results are obtained using time-series data of measles case reports from three locations during the pre-vaccine era, specifically the weekly reported cases in England and Wales (1948-1967) and monthly reported cases in New York City (1945-1964) and Baltimore (1928-1960). The proposed approach is able to well-capture the periodic, time-varying behavior of the seasonal transmission while simultaneously estimating static parameters representing the initial conditions of the system and the reporting probability of measles cases, which is vital for predicting under-reported incidence data.

2 Review: Nonlinear Filtering and the Augmented EnKF

In setting up the static parameter estimation problem, we assume an ordinary differential equation (ODE) model involving both states $x = x(t)$ and parameters θ , whose values may be uncertain or completely unknown, to describe the system dynamics, i.e.,

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

where $x = x(t) \in \mathbb{R}^d$ is the state vector, $\theta \in \mathbb{R}^p$ is the unknown (or poorly known) parameter vector, $f : \mathbb{R} \times \mathbb{R}^d \times \mathbb{R}^p \rightarrow \mathbb{R}^d$ is the known model function, and x_0 is the possibly unknown (or poorly known) initial value. We further assume discrete, noisy observations $y_k \in \mathbb{R}^m$, $k = 1, 2, \dots, T_{\text{obs}}$, of some model states,

$$y_k = g(x(t_k), \theta) + w_k, \quad 0 < t_1 < t_2 < \dots < t_{T_{\text{obs}}} \quad (2)$$

where $g : \mathbb{R}^d \times \mathbb{R}^p \rightarrow \mathbb{R}^m$, $m \leq d$, is the known observation function and w_k represents the observation error. The inverse problem is to estimate the parameter vector θ and the state vector $x(t)$ at given times from the observations y_k .

We approach the solution to the inverse problem from the Bayesian perspective, where unknowns are treated as random variables. In particular, we focus on the use of nonlinear filtering algorithms, which provide a natural setting for the time-varying parameter estimation considered in this work. In this section, we review nonlinear filtering algorithms and the augmented ensemble Kalman filter for combined state and parameter estimation, which we utilize to compute the results in Section 4.

2.1 Nonlinear Filtering Algorithms

In the Bayesian framework, the model states x and parameters θ are treated as random variables with probability distributions, and their joint posterior density is assembled using Bayes' theorem

$$\pi(x, \theta | y) \propto \pi(y | x, \theta) \pi(x, \theta) \quad (3)$$

where the likelihood function $\pi(y | x, \theta)$ indicates how likely it is that the data y are observed if the state and parameter values were known and the prior distribution $\pi(x, \theta)$ encodes any known information on the states and parameters before taking the data into account.

Filtering methods rely on the use of discrete-time stochastic equations describing the model states and observations to sequentially update the joint posterior density. Assume a time discretization t_j , $j = 0, 1, \dots, T$, with the observations in (2) occurring possibly in a subset of the discrete time instances. To avoid double indexing of the time discretization between the propagation steps and observation times, let $y_j = \emptyset$ if there is no observation at t_j . Considering both the model states and observations as Markov processes, we can write an evolution-observation model

for the stochastic state and parameter estimation problem using discrete-time Markov models. The state evolution equation

$$X_{j+1} = F(X_j, \theta) + V_{j+1}, \quad V_{j+1} \sim \mathcal{N}(0, C_{j+1}), \quad (4)$$

where F is a known propagation model and V_{j+1} is an innovation process, allows us to compute the forward time propagation of the state variables X_j given parameters θ . In this work, F computes the numerical solution to the ODE model (1) at time t_{j+1} . The observation equation

$$Y_{j+1} = G(X_{j+1}, \theta) + W_{j+1}, \quad W_{j+1} \sim \mathcal{N}(0, D_{j+1}), \quad (5)$$

where G is a known operator and W_{j+1} is the observation noise, is analogous to the observation model (2).

Denoting by D_j the accumulated observations up to time t_j ,

$$D_j = \{y_1, y_2, \dots, y_j\}, \quad (6)$$

the aim of Bayesian filtering is to sequentially update the posterior distribution $\pi(x_j, \theta \mid D_j)$ using a two-step, predictor-corrector-type scheme:

$$\pi(x_j, \theta \mid D_j) \longrightarrow \pi(x_{j+1}, \theta \mid D_j) \longrightarrow \pi(x_{j+1}, \theta \mid D_{j+1}). \quad (7)$$

The first step (known as the prediction step) uses the state evolution equation (4) to predict the values of the states at time t_{j+1} without knowledge of the data, while the second step (the analysis step or observation update) uses the observation equation (5) to correct that prediction by taking into account the data at t_{j+1} . Note that if there is no data observed at t_{j+1} , then $D_{j+1} = D_j$ and the prediction density $\pi(x_{j+1}, \theta \mid D_j)$ is equivalent to the posterior $\pi(x_{j+1}, \theta \mid D_{j+1})$. Starting with a prior density $\pi(x_0, \theta \mid D_0)$, $D_0 = \emptyset$, this updating scheme is repeated until the final joint posterior density is obtained when $j = T$.

2.2 Augmented Ensemble Kalman Filter

There are a variety of nonlinear filtering algorithms for state and parameter estimation available in the literature, including particle filters and ensemble Kalman-type filters; see, e.g., [17–25]. In this work we employ the augmented (or joint) EnKF in the style of [25], which accommodates systems of possibly very stiff differential equations. The algorithm is implemented as follows. Assume the current density $\pi(x_j, \theta \mid D_j)$ is represented in terms of an ensemble

$$\mathcal{S}_{j|j} = \left\{ (x_{j|j}^1, \theta_{j|j}^1), (x_{j|j}^2, \theta_{j|j}^2), \dots, (x_{j|j}^{N_{\text{ens}}}, \theta_{j|j}^{N_{\text{ens}}}) \right\} \quad (8)$$

where each of the N_{ens} ensemble members comprises a pair of model states $x_{j|j}$ and parameters $\theta_{j|j}$ at time t_j .

In the prediction step, the states at time t_{j+1} are predicted using the state evolution equation (4) to form a state prediction ensemble,

$$x_{j+1|j}^n = F(x_{j|j}^n, \theta_{j|j}^n) + v_{j+1}^n, \quad n = 1, 2, \dots, N_{\text{ens}}, \quad (9)$$

where $v_{j+1}^n \sim \mathcal{N}(0, C_{j+1})$ represents error in the model prediction. As in [25], we use linear multistep methods for the time integration and systematically assign the model error covariance $C_{j+1} = C_{j+1}^n$ sequentially using higher order method error control. The parameter values are not updated during the prediction step, so

$$\theta_{j+1|j}^n = \theta_{j|j}^n, \quad n = 1, 2, \dots, N_{\text{ens}}. \quad (10)$$

Prediction ensemble statistics are computed using augmented state and parameter vectors

$$z_{j+1|j}^n = \begin{bmatrix} x_{j+1|j}^n \\ \theta_{j+1|j}^n \end{bmatrix} \in \mathbb{R}^{d+k}, \quad n = 1, 2, \dots, N_{\text{ens}}, \quad (11)$$

where the prediction ensemble mean is given by

$$\bar{z}_{j+1|j} = \frac{1}{N_{\text{ens}}} \sum_{n=1}^{N_{\text{ens}}} z_{j+1|j}^n \quad (12)$$

and the prediction (or prior) covariance matrix is

$$\Gamma_{j+1|j} = \frac{1}{N_{\text{ens}} - 1} \sum_{n=1}^{N_{\text{ens}}} (z_{j+1|j}^n - \bar{z}_{j+1|j})(z_{j+1|j}^n - \bar{z}_{j+1|j})^\top. \quad (13)$$

When an observation y_{j+1} arrives, an artificial observation ensemble is generated around the true observation, such that

$$y_{j+1}^n = y_{j+1} + w_{j+1}^n, \quad n = 1, 2, \dots, N_{\text{ens}} \quad (14)$$

where $w_{j+1}^n \sim \mathcal{N}(0, D_{j+1})$ represents the observation error. The artificial observation ensemble is compared to the observation model prediction ensemble,

$$\hat{y}_{j+1}^n = g(x_{j+1|j}^n, \theta_j^n), \quad n = 1, 2, \dots, N_{\text{ens}}, \quad (15)$$

which is computed using the observation function g defined in (2). The augmented posterior ensemble at time t_{j+1} is then computed by

$$z_{j+1|j+1}^n = z_{j+1|j}^n + K_{j+1}(y_{j+1}^n - \hat{y}_{j+1}^n), \quad n = 1, 2, \dots, N_{\text{ens}} \quad (16)$$

where the Kalman gain is defined as

$$K_{j+1} = \Sigma_{j+1}^{z\hat{y}} (\Sigma_{j+1}^{\hat{y}\hat{y}} + D_{j+1})^{-1} \quad (17)$$

with $\Sigma_{j+1}^{z\hat{y}}$ denoting the cross covariance of the augmented state-parameter predictions $z_{j+1|j}^n$ and observation predictions \hat{y}_{j+1}^n , $\Sigma_{j+1}^{\hat{y}\hat{y}}$ the forecast error covariance of the observation prediction ensemble, and D_{j+1} the observation noise covariance. This formulation of the Kalman gain straightforwardly allows for nonlinear observations, as opposed to the more familiar formula for linear observation models [38]. Use of the artificial observation ensemble (14) ensures that the resulting posterior ensemble in (16) does not have too low a variance [23]. The posterior means and covariances for the states and parameters are then computed using posterior ensemble statistics, and the process repeats.

In the above treatment, the parameters θ are assumed to be static, i.e., $d\theta/dt = 0$, and are artificially evolved over time. The parameter values are not changed in the prediction step (10) and are only updated in the analysis step (16) at each data arrival through use of the cross-correlation between the parameters and model states encoded in the Kalman gain (17). If the parameter values are thought to change slowly over time, a drift can be added to the parameters in the prediction step by modeling the change in the parameter values as a random walk, thereby replacing (10) with

$$\theta_{j+1|j}^n = \theta_{j|j}^n + \xi_{j+1}^n \quad (18)$$

where $\xi_{j+1}^n \sim \mathcal{N}(0, \mathbf{E}_{j+1})$. The covariance matrix \mathbf{E}_{j+1} of the drift term must be carefully chosen a priori for each application considered in order to avoid filter divergence and obtain a useful parameter estimate. Filter divergence is a situation in which the EnKF becomes overconfident in an incorrect estimate and ignores subsequent data. This can occur either when the ensemble spread becomes too small (classical filter divergence) or too large (catastrophic filter divergence). For more details, see, e.g., [29–37].

3 Treatment of Periodic, Time-Varying Parameters

The nonlinear filtering methods and augmented EnKF reviewed in Section 2 are typically used to estimate static parameters, assuming $d\theta/dt = 0$ and artificially evolving the parameter over time to converge to a constant value. In the case of a time-varying parameter that is changing at a rate slower than the dynamics of the system, parameter tracking can be employed to estimate the change in the parameter over time, given that an appropriate drift covariance is selected to capture the drift [26–28]. However, parameter tracking does not account for known structural characteristics of the time-varying parameter throughout the estimation process. In particular, using the augmented EnKF with parameter tracking does not guarantee that the periodic variation of the parameter is fully captured or maintained in the resulting time series estimate. While it is possible to model periodic parameters using periodic functions such as sinusoids, the structural restrictions often do not adequately capture the true time evolution of the parameters [7, 8].

The goal of this work is to use instead an approach that will retain periodic structure without imposing a sinusoidal shape. To this end, we propose to model periodic, time-varying parameters as piecewise functions with unknown coefficients, repeated each period over the course of the time series, and use nonlinear filtering to estimate the coefficients. For ease of illustration, in this work we choose piecewise functions comprising a sequence of constant parameters, with each constant in the sequence representing the average value of the time-varying parameter over a specified time segment. Generally written, if $\gamma(t)$ denotes the time varying parameter with known period p , we let

$$\gamma(t_j) = \gamma_{\frac{\ell}{p} \bmod(t_j, p)} = \begin{cases} \gamma_1 & \text{if } t_j \in \left[0, \frac{p}{\ell}\right) \\ \gamma_2 & \text{if } t_j \in \left[\frac{p}{\ell}, \frac{2p}{\ell}\right) \\ \vdots & \vdots \\ \gamma_\ell & \text{if } t_j \in \left[\frac{(\ell-1)p}{\ell}, p\right) \end{cases} \quad (19)$$

where ℓ is the number of constants in the sequence. This interpretation allows us to employ nonlinear filtering (in particular, the augmented EnKF as described) to estimate the individual constant parameters γ_m , $m = 1, \dots, \ell$, for each time segment, while freeing the shape constraints on the time-varying function. This approach can be straightforwardly extended to treat $\gamma(t)$ as, e.g., a piecewise linear spline and use filtering to estimate the spline coefficients, in the same spirit of freeing the shape constraints of the resulting time-varying function. Note that the proposed approach is not restricted to the use of nonlinear filtering methods to estimate the coefficients; e.g., least squares optimization or MCMC-type sampling could be used to estimate the γ_m parameters. However, using nonlinear filtering methods allows for straightforward comparison of the proposed approach with the parameter tracking algorithms that have been utilized in this setting.

To estimate $\gamma(t)$ as formulated in (19) using the augmented EnKF, we assume that the current

density $\pi(x_j, \gamma_j \mid D_j)$ at time t_j is represented in terms of a discrete ensemble

$$\mathcal{S}_{j|j} = \left\{ (x_{j|j}^n, (\gamma_1)_{j|j}^n, \dots, (\gamma_\ell)_{j|j}^n) \right\}_{n=1}^{N_{\text{ens}}} \quad (20)$$

of size N_{ens} as in (8) and apply the algorithm as outlined in Section 2.2 to sequentially update the ℓ unknown constants γ_m comprising $\gamma(t)$.

For comparison with the proposed method, if $\gamma(t)$ changes at a rate slower than the dynamics of the system, it is possible to select an appropriate drift covariance so that the augmented EnKF with parameter tracking is able to capture the change in the parameter value over time. In this case, the discrete ensemble at time t_j is given by

$$\mathcal{S}_{j|j} = \left\{ (x_{j|j}^n, \gamma_{j|j}^n) \right\}_{n=1}^{N_{\text{ens}}}, \quad (21)$$

where each $\gamma_{j|j}^n$ represents an estimate of $\gamma(t_j)$. At the prediction step of the filter, the parameter ensemble drifts according to the random walk (18) with $\xi_{j+1}^n \sim \mathcal{N}(0, \sigma_\xi^2)$ for some prescribed variance σ_ξ^2 . We will demonstrate that while careful selection of σ_ξ allows parameter tracking to capture the time-varying behavior of the parameter, this method does not guarantee that the underlying periodicity of the parameter is maintained in the resulting time series estimate.

We note that in many applications, the parameter vector θ defined in the inverse problem (1)–(2) may generally contain a combination of both static and time-varying parameters. In this case, $\gamma(t)$ can be considered a subset of θ , and both formulations of the augmented EnKF can be straightforwardly extended to incorporate simultaneous estimation of static and time-varying parameters. This will be further demonstrated in the results.

4 Results

In this section we demonstrate the effectiveness of the proposed approach to estimating periodic, time-varying parameters on two biological examples. We first consider a synthetic example with data generated from the nonlinear FitzHugh-Nagumo system, which models the excitability of a nerve cell, to estimate the external voltage parameter, and we compare the results of the proposed method with the parameter tracking algorithm. We then perform a case study using reported measles incidence data from three locations during the pre-vaccine era to estimate the seasonal transmission parameter.

4.1 Synthetic Example: Estimating the External Voltage Parameter in the FitzHugh-Nagumo System

The FitzHugh-Nagumo system [39]

$$\frac{dx_1}{dt} = c \left(x_2 + x_1 - \frac{x_1^3}{3} + v(t) \right) \quad (22)$$

$$\frac{dx_2}{dt} = -\frac{1}{c} \left(x_1 - a + bx_2 \right) \quad (23)$$

is commonly used as a simplified version of the Hodgkin-Huxley system [40] to model the spiking dynamics of neurons. The state variable x_1 represents the measurable membrane potential of the neuron, while x_2 represents an unobservable combined effect of different ionic currents. The




Figure 1. States and voltage parameter of the FitzHugh-Nagumo system. Noisy observations of the x_1 component (top panel) of the FitzHugh-Nagumo system (22)–(23), along with the unobserved state x_2 (bottom panel, solid black) and external voltage parameter $v(t)$ (bottom panel, dashed black). In each panel, the x-axis shows time from $t = 0$ to $t = 314$ units.

parameters $a = 0.7$, $b = 0.8$, and $c = 3$ are assumed to be known and fixed, while the time-varying external voltage $v(t)$ is unknown.

We generate synthetic data by letting $v(t)$ be the negative absolute value of a cosine function, up to an additive constant, with frequency $\omega = 0.06$ that varies more slowly than the system dynamics; a similar function is used for $v(t)$ in [26]. Measurements of x_1 are taken at 943 equidistant time instances over the interval from $t = 0$ to $t = 314$, which covers three full periods of $v(t)$, as shown in Figure 1. Observations are corrupted with Gaussian noise with zero mean and standard deviation assigned to be 20% of the standard deviation of the x_1 component. Since x_2 is not observed, the observation function

$$g(x(t_k), \theta) = Gx(t_k) \quad (24)$$

is linear, with the projection matrix G picking out the x_1 component of the state vector x .

For the piecewise constant parameter estimation, we treat $v(t)$ as a sequence of 20 constant parameters v_m , $m = 1, \dots, 20$, over the course of one period, repeated across all periods as in (19), and use the augmented EnKF with $N_{\text{ens}} = 200$ ensemble members to estimate $\theta = (v_1, \dots, v_{20})$. The initial ensemble of parameter values is drawn uniformly from $\mathcal{U}(-2, 1)$. Assuming that the true initial conditions of the system are unknown, we draw the initial state ensemble for x_1 uniformly from 0.5 to 1.5 times the first observation point and let the initial state for x_2 be zero. Time integration in the prediction step of the filter is computed using Adams-Moulton methods of orders 2 and 3.

Figure 2 shows the EnKF time series estimates of the constants v_m , $m = 1, \dots, 20$, comprising the piecewise constant voltage parameter $v(t)$, and Figure 3 shows the resulting estimate of $v(t)$ using the posterior estimates of each v_m , repeated over three periods. Note that the 20 v_m parameters all converge to constant values in a sequential manner over the course of one period with very little uncertainty. The plot in Figure 3 demonstrates that using the posterior mean estimates of the v_m to define $v(t)$ as a piecewise constant function provides a fairly accurate estimate of the underlying

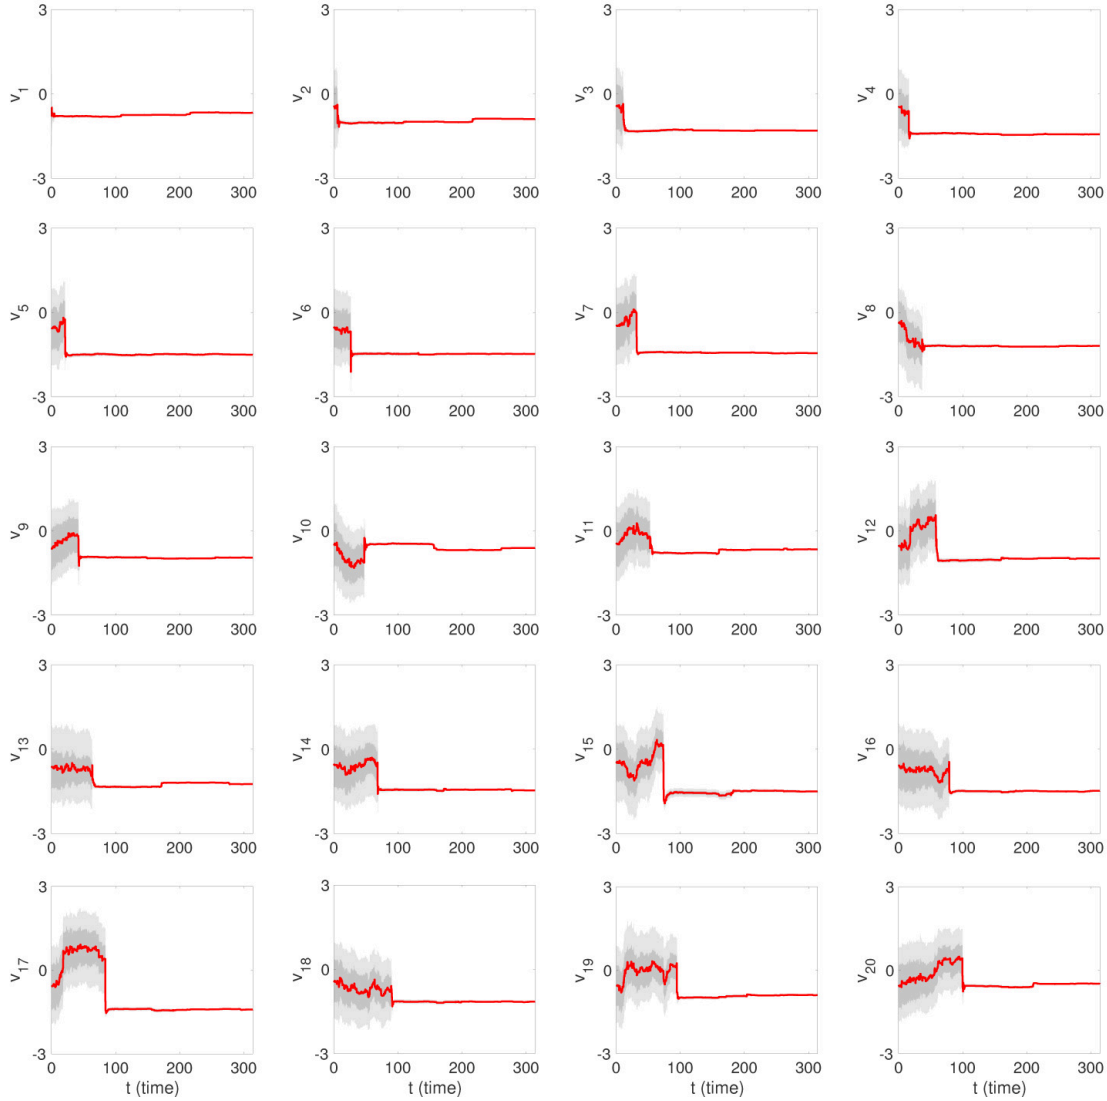


Figure 2. Parameter estimates for the piecewise constant voltage parameter. EnKF time series estimates of the constants v_m , $m = 1, \dots, 20$, comprising the piecewise constant voltage parameter $v(t)$ in the FitzHugh-Nagumo system (22)–(23). In each panel, the x-axis shows time from $t = 0$ to $t = 314$ units and the y-axis shows the value of the voltage parameter v_m . The estimated EnKF mean is plotted in solid red, and the 50% and 90% credible intervals are plotted in dark and light grey, respectively.

sinusoidal voltage function. Connecting the constant values with a linear spline provides a better visual representation of the estimated $v(t)$ curve.

For comparison, Figure 4 shows the time series estimate of $v(t)$ using the augmented EnKF with parameter tracking (18). Here $E_{j+1} = \sigma_\xi^2$ with $\sigma_\xi = 0.01$. The parameter tracking estimate of $v(t)$ well-captures the overall behavior of the voltage function over time. However, the estimate is out of phase with and does not maintain the known periodicity of the underlying function. Although not shown, both methods are able to well-recover the blind system component x_2 . The initial conditions of the system could also be estimated, as demonstrated in the next example.




Figure 3. Piecewise estimate of the voltage parameter. Posterior estimate of the piecewise constant voltage parameter $v(t)$ in the FitzHugh-Nagumo system (22)–(23), repeated over three periods. The posterior EnKF mean for each v_m is shown in solid red, connected by a linear spline. The true sinusoidal voltage function used to generate the synthetic data is plotted in dashed black.




Figure 4. Parameter tracking estimate of the voltage parameter. EnKF with parameter tracking estimate of the voltage parameter $v(t)$ in the FitzHugh-Nagumo system (22)–(23). The estimated EnKF mean is plotted in solid red, and the 50% and 90% credible intervals are plotted in dark and light grey, respectively. The true sinusoidal voltage function used to generate the synthetic data is plotted in dashed black.

4.2 Case Study: Estimating the Seasonal Transmission Parameter for Measles Incidence Data

To further demonstrate the effectiveness of the proposed methodology, we perform a case study to estimate the seasonal transmission parameter in an epidemic model for the spread of measles using reported incidence data from three locations during the pre-vaccine era. The data sets were obtained from an online infectious disease database (<http://ms.mcmaster.ca/~bolker/measdata.html>). In particular, the data comprise the weekly reported measles cases for the aggregate of 60 cities in England and Wales from 1948 to 1967, the monthly reported measles cases in New York City from 1945 to 1964, and the monthly reported measles cases in Baltimore, Maryland, from 1928 to 1960, as shown in Figure 5.

Between 1945 and the onset of widespread vaccine usage in the mid 1960's, measles outbreaks in New York City occurred about every two years in the even numbered years. Estimates show that about 1 in 8 measles cases were reported in New York City during this time [7]. Measles outbreaks in Baltimore from 1928 to 1960 were much more sporadic, occurring every two to three years with a less clear periodic pattern. Approximately 1 in 3 or 4 measles cases were reported in Baltimore during this time [7]. The measles outbreaks in England and Wales occurred more regularly, with a




Figure 5. Reported measles cases in three locations during the pre-vaccine era. The weekly reported measles cases in England and Wales (E&W) from 1948 to 1967 (left panel), monthly reported measles cases in New York City (NYC) from 1945 to 1964 (center panel), and monthly reported measles cases in Baltimore from 1928 to 1960 (right panel).

nearly annual pattern between 1948 and 1950, then following a bi-annual pattern. The reporting probability of cases in England and Wales during this time period was found to be relatively high, with greater than 50% of cases reported [41]. For more details on the data, see [7, 42–44].

We use a four compartment Susceptible-Exposed-Infectious-Recovered (SEIR) model to predict the epidemic system dynamics in this application. For a review of SIR-type models in epidemiology, see, e.g., [45, 46]. A standard SEIR model comprises the following system of ODEs:

$$\frac{dS}{dt} = m(N - S) - \frac{\beta(t)SI}{N} \quad (25)$$

$$\frac{dE}{dt} = \frac{\beta(t)SI}{N} - (m + a)E \quad (26)$$

$$\frac{dI}{dt} = aE - (m + c)I \quad (27)$$

$$\frac{dR}{dt} = cI - mR \quad (28)$$

where m is the birth rate of new susceptible individuals and the death rate of individuals in each compartment (assumed here to be equal), a is the per-capita rate at which exposed individuals become infectious, c is the per-capita recovery rate of infectious individuals, and $\beta(t)$ is the unknown seasonal transmission parameter, which varies with time over the course of a year. Assuming a constant population size N , the system (25)–(28) can be reduced to

$$\frac{dS}{dt} = m(N - S) - \frac{\beta(t)SI}{N} \quad (29)$$

$$\frac{dE}{dt} = \frac{\beta(t)SI}{N} - (m + a)E \quad (30)$$

$$\frac{dI}{dt} = aE - (m + c)I \quad (31)$$

where

$$R(t) = N - S(t) - E(t) - I(t). \quad (32)$$

As described above, the data are the recorded number of measles cases reported over a specified time period (e.g., weekly, monthly). Since measles cases are known to be under-reported [47,

48], each observation $y_k = y(t_k)$ is modeled as a reported fraction of the total number of cases accumulated between times t_{k-1} and t_k . The observation model is formulated as in (2) with nonlinear observation function

$$g(x(t_k), \theta) = \rho d(x(t_k), \theta) \quad (33)$$

where ρ is the reporting probability, assumed here to be constant over time, and

$$d(x(t_k), \theta) = \int_{t_{k-1}}^{t_k} \frac{\beta(t)S(t)I(t)}{N} dt \quad (34)$$

denotes the total number of cases between times t_{k-1} and t_k . The observation function in (33) can be written equivalently as

$$g(x(t_k), \theta) = \int_{t_{k-1}}^{t_k} \rho \frac{\beta(t)S(t)I(t)}{N} dt \quad (35)$$

where

$$\frac{dg}{dt} = \rho \frac{\beta(t)S(t)I(t)}{N} \quad (36)$$

gives the rate of change of the cumulative cases between times t_{k-1} and t_k .

We impose an annually-varying, periodic structure on the seasonal transmission parameter by modeling $\beta(t)$ as a sequence of constant parameters β_m , $m = 1, \dots, 12$, with each β_m representing the transmission parameter for a given month, as in (19); i.e.,

$$\beta(t_j) = \beta_{\text{mod}(t_j, 12)} = \begin{cases} \beta_1 & \text{if } t_j \in \text{January} \\ \beta_2 & \text{if } t_j \in \text{February} \\ \vdots & \vdots \\ \beta_{12} & \text{if } t_j \in \text{December} \end{cases} \quad (37)$$

Since we assume that the same seasonal pattern is repeated annually [43], the period of the transmission parameter is one year.

In addition to $\beta(t)$, there are several static system parameters that need to be estimated in this application. While the birth/death rate m , the exposed-to-infectious rate a , the recovery rate c , and the population size N in (29)–(31) can be fairly well estimated from the literature and demographic data, the initial model states $S(0)$, $E(0)$, and $I(0)$ remain uncertain. Further, while some estimates of the reporting probability ρ can be obtained [47, 48], this parameter is also uncertain. Therefore, the unknown parameter vector to be estimated is $\theta = (\beta_1, \dots, \beta_{12}, S(0), E(0), I(0), \rho) \in \mathbb{R}^{16}$.

Synthetic validation was performed on the measles application prior to using the real data; results obtained using synthetic data are provided as supplementary material in the Appendix. Results using the real data are obtained using a similar procedure to the synthetic examples, incorporating the details specific to each data set in the filter setup. In each simulation, the regional population sizes are assumed to be constant and are approximated from available demographic data (see <http://ms.mcmaster.ca/~bolker/measdata.html>). In particular, we assume the population size for England and Wales is approximately 40,000,000, New York City is 7,800,000, and Baltimore is 891,080, during the respective time spans over which the measles cases were recorded.

Prior ensembles of the initial model states are drawn by assuming that the population at each location initially comprises 5% susceptible individuals, 0.8% exposed, and 0.2% infected, then sampling from a uniform distributions between 0.8 and 1.2 times those values. Prior ensembles for the reporting probabilities are drawn using available reporting information for each location during




Figure 6. Piecewise estimates of the transmission parameter in each location. Posterior estimates of the piecewise constant transmission parameters $\beta(t)$ corresponding to the reported measles cases in England and Wales (E&W), New York City (NYC), and Baltimore. In each figure, the posterior EnKF mean for each constant β_m , corresponding to each month in a year, is shown in solid red, connected by a linear spline. Corresponding posterior estimates of the static parameters $S(0)$, $E(0)$, $I(0)$, and ρ for each location are listed in Table 1.

the years considered [7, 41]. In particular, for England and Wales the initial distribution of ρ is drawn uniformly between 0.55 and 0.75, for New York City between 0.05 and 0.2, and for Baltimore between 0.2 and 0.4.

Figure 6 shows the resulting piecewise constant estimate of $\beta(t)$ for the three respective locations, using the posterior estimates of each of the 12 estimated β_m constants for each data set, connected with linear splines. Table 1 lists the corresponding static parameter estimates for $S(0)$, $E(0)$, $I(0)$, and ρ for each location. The resulting shapes of the transmission parameters reflect the annual pattern that we might expect, with the lowest value of $\beta(t)$ in each location occurring during the summer months, corresponding to summer holiday break for school-aged children.

As the summer break in the United Kingdom typically begins later than in the United States, the minimum value of the transmission parameter shown in Figure 6 for the England and Wales data occurs a bit later (i.e., in September) than in the US cities (July for New York City, August

	E&W	NYC	Baltimore
$S(0)$	1,531,000	569,300	4,1250
$E(0)$	245,000	91,080	6,600
$I(0)$	61,240	22,770	1,650
ρ	0.7316	0.1282	0.3270

Table 1. Static parameter estimates for the reported measles cases in each of three locations. Posterior estimates of the static parameters $S(0)$, $E(0)$, $I(0)$, and ρ for the reported measles cases in England and Wales (E&W), New York City (NYC), and Baltimore, respectively, which correspond to the estimated transmission parameters $\beta(t)$ shown in Figure 6. Parameter values are listed with four significant figures.

for Baltimore). The transmission parameter noticeably rises during the fall months (September, October) when children return to school, before decreasing again during the winter holiday break. There is also a clear dip in $\beta(t)$ during the spring (March, April), which may coincide with a mid-semester break. The gradual drop seen in $\beta(t)$ during certain times over the course of the semester may be the result of imperfect mixing; see, e.g., [43, 44, 49].

5 Discussion

In this work we present an approach to estimating periodic, time-varying forcing parameters in nonlinear systems through use of nonlinear filtering methodology. We demonstrate the effectiveness of the approach using two applications from the life sciences, namely estimating the external voltage parameter in the FitzHugh-Nagumo system for neuronal spiking and the seasonal transmission parameter in an epidemic model for the spread of measles. By treating the forcing parameters in these applications as piecewise constant functions and using nonlinear filtering methodology to estimate their coefficients, we are able to incorporate the periodic structure of the parameters into the estimation process without prescribing restrictive evolution models. We are also able to simultaneously estimate time-invariant parameters associated with the system, which for the measles example includes the initial conditions of the model states and the reporting probability of cases for the data considered.

The approach presented in this work is meant as a step towards better estimating time-varying parameters by incorporating known structural characteristics, in particular, periodicity. For the applications at hand, treating the forcing parameter as a periodically repeating piecewise constant function is a middle ground approach between assigning a strict functional shape and letting the parameter drift with no guaranteed structure. Assigning, e.g., a sinusoidal function does not necessarily represent the true behavior of the time-varying parameter of interest for applications with real data. This can be seen in the results in Figure 6 for the measles data, where the seasonal transmission parameters corresponding to the real data sets analyzed in this work noticeably do not follow strict sinusoidal patterns over the course of a year.

While the piecewise constant approach requires that ℓ constant parameters be estimated, it provides a way to enforce periodicity over the full time series of the data without imposing a restrictive functional form. The parameter tracking approach, while requiring that fewer parameters be estimated during the filtering process (i.e., one dynamic parameter vs. ℓ constants), introduces an additional nuisance parameter (namely, the drift variance σ_ξ^2) which must be carefully chosen a priori for each application considered in order to avoid filter divergence and obtain a useful estimate. The parameter tracking approach also does not guarantee the resulting time-series estimate will maintain the periodic structure of the parameter over the full time series, as seen with the FitzHugh-Nagumo external voltage parameter in Figure 4. Future work may include combining these two approaches to better incorporate known structural characteristics into a parameter tracking-type filtering algorithm to allow for more freedom in the shape of the resulting time-varying parameter estimate.

It is indeed possible to use a variety of parameter estimation techniques, such as least squares optimization or MCMC sampling, to estimate the ℓ constant coefficients in the piecewise constant approach. In this work, we highlight the use of nonlinear filtering algorithms, specifically the augmented EnKF, since these methods are amenable to the time-series data often available in life sciences applications and the resulting posterior distributions contain a measure of uncertainty in the parameter estimates. While not the focus of this work, the posterior parameter distributions can be used for model prediction and uncertainty quantification, which would require additional

variational techniques in the deterministic setting [12, 50]. Further, using nonlinear filtering to estimate the coefficients in the proposed method allows for straightforward comparison with the parameter tracking algorithm as described. While we use the augmented EnKF to obtain our results, a variety of other sequential filtering algorithms could also be employed, e.g., a particle filter [17–21] or a dual filter, either heuristic [38] or Bayesian-consistent [51].

As previously noted, the proposed method is not restricted to the use of piecewise constant functions and can be straightforwardly extended to treat the time-varying parameter as, e.g., a piecewise linear spline and use nonlinear filtering to estimate the spline coefficients. The piecewise constant interpretation used in this work is meant as a simple example of a more general framework for estimating periodic, time-varying parameters in nonlinear systems. While in this work we consider only one time-varying parameter per system to estimate, additional model parameters, such as the reporting probability ρ in the measles application, could be treated as time-varying, assuming some known structural characteristics and taking into account the added computational cost.

Acknowledgments

This work was supported by National Science Foundation grant RTG/DMS-1246991 (Research Training Group in Mathematical Biology at North Carolina State University).

References

[1] S Altizer, A Dobson, P Hosseini, P Hudson, M Pascual, and P Rohani. Seasonality and the dynamics of infectious diseases. *Ecology Letters*, 9:467–484, 2006.

[2] J L Aron and I B Schwartz. Seasonality and period-doubling bifurcations in an epidemic model. *J Theor Biol*, 110:665–679, 1984.

[3] N C Grassly and C Fraser. Seasonal infectious disease epidemiology. *Proc R Soc B*, 273:2541–2550, 2006.

[4] C B Saper, T E Scammell, and J Lu. Hypothalamic regulation of sleep and circadian rhythms. *Nature*, 437:1257–1263, 2005.

[5] J-C Leloup and A Goldbeter. Toward a detailed computational model for the mammalian circadian clock. *PNAS*, 100(12):7051–7056, 2003.

[6] R E Kronauer, C A Czeisler, S F Pilato, M C Moore-Ede, and E D Weitzman. Mathematical model of the human circadian system with two interacting oscillators. *American Journal of Physiology: Regulatory, Integrative and Comparative Physiology*, 242(1):R3–R17, 1982.

[7] W P London and J A Yorke. Recurrent outbreaks of measles, chickenpox and mumps: I, seasonal variation in contact rates. *Am. J. Epidemiol.*, 98:453–468, 1973.

[8] K Dietz. The incidence of infectious diseases under the influence of seasonal fluctuations. *Lect. Notes Biomath.*, 11:1–15, 1976.

[9] D W Marquardt. An algorithm for least-squares estimation of nonlinear parameters. *J Soc Indust Appl Math*, 11(2):431–441, 1963.

- [10] M L Johnson and L M Faunt. Parameter estimation by least-squares methods. *Methods in Enzymology*, 210:1–37, 1992.
- [11] J E Dennis Jr. and R B Schnabel. *Numerical Methods for Unconstrained Optimization and Nonlinear Equations*. Classics in Applied Mathematics. SIAM, Philadelphia, 1996.
- [12] H T Banks, S Hu, and W C Thompson. *Modeling and Inverse Problems in the Presence of Uncertainty*. CRC Press, New York, 2014.
- [13] C Andrieu and J Thoms. A tutorial on adaptive MCMC. *Statistics and Computing*, 18(4):343–373, 2008.
- [14] H Haario, E Saksman, and J Tamminen. Adaptive proposal distribution for random walk Metropolis algorithm. *Comput Statist*, 14:375–395, 1999.
- [15] H Haario, E Saksman, and J Tamminen. An adaptive Metropolis algorithm. *Bernoulli*, 7:223–242, 2001.
- [16] H Haario, M Laine, A Mira, and E Saksman. DRAM: Efficient adaptive MCMC. *Statistics and Computing*, 16:339–354, 2006.
- [17] J P Kaipio and E Somersalo. *Statistical and Computational Inverse Problems*. Applied Mathematical Sciences. Springer, New York, 2005.
- [18] J Liu and M West. Combined parameter and state estimation in simulation-based filtering. In A Doucet, N de Freitas, and N Gordon, editors, *Sequential Monte Carlo Methods in Practice*, pages 197–223, New York, 2001. Springer.
- [19] M Pitt and N Shephard. Filtering via simulation: auxiliary particle filters. *J Amer Statist Assoc*, 94:590–599, 1999.
- [20] E L Ionides, C Breto, and A A King. Inference for nonlinear dynamical systems. *PNAS*, 103(49):18438–18443, 2006.
- [21] A Arnold, D Calvetti, and E Somersalo. Linear multistep methods, particle filtering and sequential Monte Carlo. *Inverse Problems*, 29(8):085007, 2013.
- [22] G Evensen. 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.
- [23] G Burgers, P J van Leeuwen, and G Evensen. Analysis scheme in the ensemble Kalman filter. *Mon Weather Rev*, 126(6):1719–1724, 1998.
- [24] G Evensen. The ensemble Kalman filter for combined state and parameter estimation. *IEEE Control Syst Mag*, 29(3):83–104, 2009.
- [25] A Arnold, D Calvetti, and E Somersalo. Parameter estimation for stiff deterministic dynamical systems via ensemble Kalman filter. *Inverse Problems*, 30(10):105008, 2014.
- [26] H U Voss, J Timmer, and J Kurths. Nonlinear dynamical system identification from uncertain and indirect measurements. *International Journal of Bifurcation and Chaos*, 14(6):1905–1933, 2004.

[27] F Hamilton, T Berry, N Peixoto, and T Sauer. Real-time tracking of neuronal network structure using data assimilation. *Physical Review E*, 88:052715, 2013.

[28] B J Matzuka. *Nonlinear Filtering Methodologies for Parameter Estimation and Uncertainty Quantification in Noisy, Complex Biological Systems*. PhD thesis, North Carolina State University, 2014.

[29] A H Jazwinski. *Stochastic Processes and Filtering Theory*. Academic Press, San Diego, CA, 1970.

[30] P L Houtekamer and H L Mitchell. Data assimilation using an ensemble Kalman filter technique. *Mon Weather Rev*, 126:796–811, 1998.

[31] T M Hamill, J S Whitaker, and C Snyder. Distance-dependent filtering of background error covariance estimates in an ensemble Kalman filter. *Mon Weather Rev*, 129:2776–2790, 2001.

[32] J L Anderson. An ensemble adjustment Kalman filter for data assimilation. *Mon Weather Rev*, 129:2884–2903, 2001.

[33] J S Whitaker and T M Hamill. Ensemble data assimilation without perturbed observations. *Mon Weather Rev*, 130:1913–1924, 2002.

[34] G-H C Ng, D McLaughlin, D Entekhabi, and A Ahanin. The role of model dynamics in ensemble Kalman filter performance for chaotic systems. *Tellus A*, 63(5):958–977, 2011.

[35] T Berry and T Sauer. Adaptive ensemble Kalman filtering of non-linear systems. *Tellus A: Dynamic Meteorology and Oceanography*, 65(1):20331, 2013.

[36] J Harlim and A J Majda. Catastrophic filter divergence in filtering nonlinear dissipative systems. *Commun Math Sci*, 8(1):27–43, 2010.

[37] G A Gottwald and A J Majda. A mechanism for catastrophic filter divergence in data assimilation for sparse observation networks. *Nonlin Processes Geophys*, 20:705–712, 2013.

[38] H Moradkhani, S Sorooshian, H V Gupta, and P R Houser. Dual state-parameter estimation of hydrological models using ensemble Kalman filter. *Adv Water Resour*, 28(2):135–147, 2005.

[39] R FitzHugh. Impulses and physiological states in theoretical models of nerve membrane. *Biophys J*, 1:445–466, 1961.

[40] A L Hodgkin and A F Huxley. A quantitative description of membrane current and its application to conduction and excitation in nerve. *J Physiol*, 117:500–544, 1952.

[41] J A Clarkson and P E M Fine. The efficiency of measles and pertussis notification in England and Wales. *International Journal of Epidemiology*, 14:153–168, 1985.

[42] P E M Fine and J A Clarkson. Measles in England and Wales – I: An analysis of factors underlying seasonal patterns. *International Journal of Epidemiology*, 11(1):5–14, 1982.

[43] B F Finkenstädt and B T Grenfell. Time series modelling of childhood diseases: a dynamical systems approach. *Applied Statistics*, 49:187–205, 2000.

[44] O N Bjørnstad, B F Finkenstädt, and B T Grenfell. Dynamics of measles epidemics: estimating scaling of transmission rates using a time series SIR model. *Ecological Monographs*, 72(2):169–184, 2002.

- [45] R M Anderson and R M May. *Infectious Diseases of Humans: Dynamics and Control*. Oxford, 1992.
- [46] H W Hethcote. The mathematics of infectious diseases. *SIAM Review*, 42(4):599–653, 2000.
- [47] D He, E L Ionides, and A A King. Plug-and-play inference for disease dynamics: measles in large and small populations as a case study. *J R Soc Interface*, pages 1–13, 2009.
- [48] C E Gunning, E Erhardt, and H J Wearing. Conserved patterns of incomplete reporting in pre-vaccine era childhood diseases. *Proc R Soc B*, 281:20140886, 2014.
- [49] B T Grenfell, O N Bjørnstad, and B F Finkenstädt. Dynamics of measles epidemics: scaling noise, determinism, and predictability with the TSIR model. *Ecological Monographs*, 72(2):185–202, 2002.
- [50] R C Smith. *Uncertainty Quantification: Theory, Implementation, and Applications*. SIAM, 2013.
- [51] B Ait-El-Fquih, M El Gharamti, and I Hoteit. A Bayesian consistent dual ensemble Kalman filter for state-parameter estimation in subsurface hydrology. *Hydrol Earth Syst Sci*, 20(8):3289–3307, 2016.
- [52] J A Nelder and R Mead. A simplex method for function minimization. *The Computer Journal*, 7:308, 1965.
- [53] J E Dennis Jr. and D J Woods. Optimization on microcomputers: the Nelder-Mead simplex algorithm. In A Wouk, editor, *New Computing Environments: Microcomputers in Large-scale Computing*, pages 116–122, Philadelphia, 1987. SIAM.

Appendix

Validation of the Method on Synthetic Measles Incidence Data

Prior to using the real measles data described in the manuscript, we performed validation of the proposed method for estimating periodic, time-varying parameters using synthetic incidence data. The inverse problem considered is to track the model states $x = (S, E, I) \in \mathbb{R}^3$ of the reduced SEIR model (29)–(31) and estimate the time-varying seasonal transmission parameter $\beta(t)$, along with the static model initial conditions $S(0)$, $E(0)$, and $I(0)$ and reporting probability ρ , given monthly data on the reported number of measles cases, with the observation function modeled as in (35). For the problem at hand, we assume that the parameters N , m , a and c in system (29)–(31) are known and fixed, so the unknown parameter vector that we want to estimate is defined as $\theta = (\beta_1, \dots, \beta_{12}, S(0), E(0), I(0), \rho) \in \mathbb{R}^{16}$.

We generate synthetic data using the fixed parameters

$$N = 9.235 \times 10^6, \quad m = 0.02, \quad a = 35.84, \quad c = 100 \quad (\text{A1})$$

and the sinusoidal transmission function

$$\hat{\beta}(t) = b_0 \left(1 + b_1 \cos(2\pi t) \right) \quad (\text{A2})$$




Figure A1. Synthetic measles incidence data. Synthetic monthly incidence data generated using the sinusoidal transmission function (A2) with average transmission $b_0 = 1800$ and amplitude b_1 chosen to simulate low seasonality ($b_1 = 0.08$) and high seasonality ($b_1 = 0.3$).

with average transmission $b_0 = 1800$ and amplitude of variation b_1 , which can be viewed as the first term of a Fourier expansion of the underlying seasonal function. Figure A1 shows two synthetic data sets, generated to simulate low seasonality (with $b_1 = 0.08$) and high seasonality (with $b_1 = 0.3$), respectively. In each case, we simulate the model states for 120 years, then discard the transient and record the total number of measles cases per month over 10 years. We assume that 60% of cases are reported and that the observations are corrupted by a small amount of Gaussian noise with standard deviation $\sigma = 0.1$. For demonstration purposes, we show results relating to the low seasonality data; although not shown, similar results are obtained using the high seasonality data.

To validate the proposed approach, in which we treat $\beta(t)$ as a sequence of constant parameters β_m , $m = 1, \dots, 12$, with each β_m representing the transmission parameter for a given month, i.e.,

$$\beta(t_j) = \beta_{\text{mod}(t_j, 12)} = \begin{cases} \beta_1 & \text{if } t_j \in \text{January} \\ \beta_2 & \text{if } t_j \in \text{February} \\ \vdots & \vdots \\ \beta_{12} & \text{if } t_j \in \text{December} \end{cases} \quad (\text{A3})$$

as defined in (37), we use the augmented EnKF outlined in Section 2.2 with $N_{\text{ens}} = 250$ ensemble members to estimate the parameters $\theta = (\beta_1, \dots, \beta_{12}, S(0), E(0), I(0), \rho)$ as described. We draw the initial parameter ensemble, for $n = 1, \dots, N_{\text{ens}}$, using uniform distributions as follows: $(\beta_m)_{0|0}^n \sim \mathcal{U}(1000, 2500)$ for $m = 1, \dots, 12$, $x_{0|0}^n \sim \mathcal{U}(0.25\hat{x}_0, 2\hat{x}_0)$ where $\hat{x}_0 = (\hat{S}(0), \hat{E}(0), \hat{I}(0)) \in \mathbb{R}^3$ is the initial condition of the states used to generate the data, and $\rho_{0|0}^n \sim \mathcal{U}(0.5, 0.75)$. Time integration in the prediction step of the filter is computed using Adams-Moulton methods of orders 1 and 2.

Figure A2 shows the EnKF time series estimates of the constants β_m , $m = 1, \dots, 12$, comprising the piecewise constant transmission parameter $\beta(t)$ defined in (A3) for the synthetic data with low seasonality. Figure A3 shows the resulting estimate of $\beta(t)$ using the posterior estimates of each β_m , repeated over 10 years, along with the corresponding time series estimates of the static parameters. Note that the 12 β_m parameters and four additional static parameters all converge to constant values around year 6 with very little uncertainty. The top panel in Figure A3 demonstrates that using the posterior mean estimates of the β_m to define the piecewise constant $\beta(t)$ provides a fairly accurate estimate of the true sinusoidal transmission function $\hat{\beta}(t)$ used to generate the data. Connecting the piecewise constant values with a linear spline provides a better visual representation of the estimated $\beta(t)$ curve.

As previously noted, the proposed approach for estimating periodic, time-varying parameters is not restricted to use of the augmented EnKF; various parameter estimation techniques could be utilized to estimate the β_m coefficients and additional static parameters. Table A1 compares the mean estimates of β_m , $m = 1, \dots, 12$, and the static parameters $S(0)$, $E(0)$, $I(0)$, and ρ computed using the augmented EnKF formulation (shown in Figures A2 and A3) with estimates obtained using a least squares optimization routine with a random initial guess. The optimization was performed utilizing MATLAB's `fminsearch` function, which employs a Nelder-Mead simplex method for unconstrained nonlinear minimization [52, 53], to find the parameter values that minimize the sum of squared errors between the observed case data and the observation model defined in (35). The initial guess for θ was chosen similarly to the augmented EnKF prior ensemble, with random draws $(\beta_m)_0 \sim \mathcal{U}(1000, 2500)$ for $m = 1, \dots, 12$, $x_0 \sim \mathcal{U}(0.25\hat{x}_0, 2\hat{x}_0)$, and $\rho_0 \sim \mathcal{U}(0.5, 0.75)$. The optimizer was set to compute a maximum of 20,000 function evaluations.

The “true” values of the β_m coefficients listed in Table A1 are the values of the true sinusoidal function $\hat{\beta}(t)$ in (A2) evaluated at the midpoint of each month over the course of a year (in units of years). The relative error in each case is computed as

$$\left| \frac{\theta_{\text{true}} - \theta_{\text{est}}}{\theta_{\text{true}}} \right| \quad (\text{A4})$$

where θ_{est} is the estimated parameter vector (either by augmented EnKF or least squares) and




Figure A2. Parameter estimates for the piecewise constant transmission parameter for the low seasonality synthetic data. EnKF time series estimates of the constants β_m , $m = 1, \dots, 12$, comprising the piecewise constant transmission parameter $\beta(t)$ defined in (A3) for synthetic data with low seasonality. In each panel, the x-axis shows time from 0 to 10 years and the y-axis shows the value of the transmission parameter β_m . The estimated EnKF mean is plotted in solid red, and the 50% and 90% credible intervals are plotted in dark and light grey, respectively.




Figure A3. Piecewise estimate of the transmission parameter, with corresponding static parameters, for the low seasonality synthetic data. (Top panel) Posterior estimate of the piecewise constant transmission parameter $\beta(t)$ defined in (A3), repeated over 10 years, for the low seasonality synthetic data. The posterior EnKF mean for each β_m , corresponding to each month in a year, is shown in solid red, connected by a linear spline. The true sinusoidal transmission function $\hat{\beta}(t)$ in (A2) used to generate the low seasonality synthetic data is plotted in dashed black. (Bottom panel) EnKF time series estimates of the static parameters $S(0)$, $E(0)$, $I(0)$, and ρ , estimated simultaneously along with β_m , $m = 1, \dots, 12$. The estimated EnKF mean is plotted in solid red, and the 50% and 90% credible intervals are plotted in dark and light grey, respectively. The true parameter values used to generate the low seasonality synthetic data are plotted in dashed black.

θ_{true} is the true parameter vector. The results in Table A1 show that the parameter estimates are fairly similar in both cases, with comparable relative errors. While not the focus of this work, it is possible to analyze the computational advantages of using nonlinear filtering in this setting, as these algorithms require time integration only from one point to the next, instead of integrating over the full time series of data at each update of the parameter estimate.

For comparison with the proposed approach, Figure A4 shows the time series estimate of $\beta(t)$ for the low seasonality data using the augmented EnKF with parameter tracking, along with the corresponding time series estimates of the four static parameters. Here, the covariance of the drift term is set as $E_{j+1} = \sigma_\xi^2$ with $\sigma_\xi = 3.8$. After around year 7, the parameter tracking estimate of $\beta(t)$ is able to capture the overall behavior of the transmission. However, the estimate is out of phase with the true underlying transmission and the resulting shape does not fully capture the annual seasonal variation over the entire time series. Further, the corresponding posterior estimates of the static parameters $S(0)$, $E(0)$, $I(0)$, and ρ are not as accurate as in the piecewise case.

As previously noted, the covariance matrix E_{j+1} of the drift term in the parameter tracking algorithm must be chosen carefully in order to avoid filter divergence and obtain a useful parameter estimate. The parameter tracking estimate for this example is sensitive to the choice of the drift variance σ_ξ^2 . To demonstrate this, Figure A5 shows the resulting parameter tracking estimate of $\beta(t)$ for the low seasonality data with $\sigma_\xi = 0.5$. Note that in this case the parameter estimate diverges around year 3, which is mirrored in the tracking of the susceptible population $S(t)$ (not

shown) and also affects the estimates of the static parameters $S(0)$, $E(0)$, $I(0)$, and ρ . This signifies the importance of carefully selecting σ_ξ in the parameter tracking scheme for this example in order to obtain a useful result.

	True	EnKF Mean	Rel. Error ($\times 10^{-3}$)	Least Squares	Rel. Error ($\times 10^{-3}$)
β_1	1,939.1	1,940.7	0.8485	1,943.8	2.4435
β_2	1,901.8	1,904.1	1.2222	1,899.9	1.0004
β_3	1,837.3	1,838.7	0.7908	1,839.4	1.1612
β_4	1,762.7	1,766.0	1.8449	1,763.6	0.4718
β_5	1,698.2	1,696.7	0.8766	1,698.0	0.1307
β_6	1,660.9	1,659.6	0.7770	1,660.2	0.4054
β_7	1,660.9	1,658.9	1.1906	1,659.0	1.1514
β_8	1,698.2	1,695.9	1.3513	1,698.0	0.0951
β_9	1,762.7	1,761.7	0.5697	1,762.5	0.1579
β_{10}	1,837.3	1,834.3	1.6356	1,837.9	0.3245
β_{11}	1,901.8	1,902.8	0.5101	1,903.3	0.7933
β_{12}	1,939.1	1,940.6	0.7678	1,940.1	0.5110
$S(0)$	553,020.0	553,730.0	1.2727	552,370.0	1.1752
$E(0)$	8,042.9	8,053.1	1.2727	8,832.2	98.1410
$I(0)$	2,765.1	2,768.7	1.2727	2,058.7	255.4700
ρ	0.6000	0.6001	0.1230	0.6002	0.3096

Table A1. Parameter estimates: augmented EnKF vs. least squares optimization.

Estimates of the 12 β_m parameters comprising the piecewise constant formulation of $\beta(t)$ and the four additional static parameters $S(0)$, $E(0)$, $I(0)$, and ρ obtained using the augmented EnKF and least squares optimization via MATLAB's `fminsearch`. The true parameter values are listed along with the posterior mean estimates obtained using the augmented EnKF and the least squares estimates. The relative errors as defined in (A4) are given for both methods. Computed values are listed with five significant figures.




Figure A4. Parameter tracking estimate of the transmission parameter, with corresponding static parameters, for the low seasonality synthetic data. (Top panel) EnKF with parameter tracking estimate of the transmission parameter $\beta(t)$ for the low seasonality synthetic data with $\sigma_{\xi} = 3.8$. The estimated EnKF mean is plotted in solid red, and the 50% and 90% credible intervals are plotted in dark and light grey, respectively. The true sinusoidal transmission function $\hat{\beta}(t)$ in (A2) used to generate the low seasonality synthetic data is plotted in dashed black. (Bottom panel) EnKF time series estimates of the static parameters $S(0)$, $E(0)$, $I(0)$, and ρ , estimated simultaneously along with $\beta(t)$. The estimated EnKF mean is plotted in solid red, and the 50% and 90% credible intervals are plotted in dark and light grey, respectively. The true parameter values used to generate the low seasonality synthetic data are plotted in dashed black.




Figure A5. Diverging parameter tracking estimate of the transmission parameter, with corresponding static parameters, for the low seasonality synthetic data. (Top panel) EnKF with parameter tracking estimate of the transmission parameter $\beta(t)$ for the low seasonality synthetic data with $\sigma_\xi = 0.5$. The estimated EnKF mean is plotted in solid red, and the 50% and 90% credible intervals are plotted in dark and light grey, respectively. The true sinusoidal transmission function $\hat{\beta}(t)$ in (A2) used to generate the low seasonality synthetic data is plotted in dashed black. (Bottom panel) EnKF time series estimates of the static parameters $S(0)$, $E(0)$, $I(0)$, and ρ , estimated simultaneously along with $\beta(t)$. The estimated EnKF mean is plotted in solid red, and the 50% and 90% credible intervals are plotted in dark and light grey, respectively. The true parameter values used to generate the low seasonality synthetic data are plotted in dashed black.