



**HAL**  
open science

# Inference for the stochastic FitzHugh-Nagumo model from real action potential data via approximate Bayesian computation

Adeline Samson, Massimiliano Tamborrino, Irene Tubikanec

## ► To cite this version:

Adeline Samson, Massimiliano Tamborrino, Irene Tubikanec. Inference for the stochastic FitzHugh-Nagumo model from real action potential data via approximate Bayesian computation. Computational Statistics and Data Analysis, In press. hal-04600178

**HAL Id: hal-04600178**

**<https://hal.science/hal-04600178v1>**

Submitted on 4 Jun 2024

**HAL** is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

# Inference for the stochastic FitzHugh-Nagumo model from real action potential data via approximate Bayesian computation

Adeline Samson\*, Massimiliano Tamborrino†, Irene Tubikanec‡

## Abstract

The stochastic FitzHugh-Nagumo (FHN) model considered here is a two-dimensional nonlinear stochastic differential equation with additive degenerate noise, whose first component, the only one observed, describes the membrane voltage evolution of a single neuron. Due to its low-dimensionality, its analytical and numerical tractability, and its neuronal interpretation, it has been used as a case study to test the performance of different statistical methods in estimating the underlying model parameters. Existing methods, however, often require complete observations, non-degeneracy of the noise or a complex architecture (e.g., to estimate the transition density of the process, “recovering” the unobserved second component), and they may not (satisfactorily) estimate all model parameters simultaneously. Moreover, these studies lack real data applications for the stochastic FHN model. Here, we tackle all challenges (non-globally Lipschitz drift, non-explicit solution, lack of available transition density, degeneracy of the noise, and partial observations) via an intuitive and easy-to-implement sequential Monte Carlo approximate Bayesian computation algorithm. The proposed method relies on a recent computationally efficient and structure-preserving numerical splitting scheme for synthetic data generation, and on summary statistics exploiting the structural properties of the process. We succeed in estimating all model parameters from simulated data and, more remarkably, real action potential data of rats. The presented novel real-data fit may broaden the scope and credibility of this classic and widely used neuronal model.

## Keywords

FitzHugh-Nagumo model, Action potential data, Hypoellipticity, Splitting numerical methods, Bayesian inference, Sequential Monte Carlo approximate Bayesian computation.

## AMS subject classifications

60H10, 60H35, 65C30

## Acknowledgements

A part of this paper was written while I.T. was member of the Department of Statistics, University of Klagenfurt, 9020 Klagenfurt, Austria. She is thankful for the support and the high-performance computing infrastructure provided by that university.

---

\*Laboratoire Jean Kuntzmann, University Grenoble Alpes (Adeline.Leclercq-Samson@univ-grenoble-alpes.fr)

†Department of Statistics, University of Warwick (massimiliano.tamborrino@warwick.ac.uk)

‡Institute of Applied Statistics, Johannes Kepler University Linz (irene.tubikanec@jku.at)

# 1 Introduction

The FitzHugh-Nagumo (FHN) model [16, 42] is a well-known mathematical model, which is applied in the fields of theoretical and computational neuroscience and used to describe the pulses and bursts occurring in the membrane potential of a single neuron. The FHN model is a nonlinear second-order polynomial differential system bearing two states representing the membrane potential  $V$  and a recovery variable  $U$ , respectively, with the latter comprising all ion channel related processes into one state. This system is less complex than the Hodgkin-Huxley [21] or the Morris-Lecar [41] model, since it does not use physical parameters like ion conductances. Stochastic versions of these neuronal models have been proposed to describe various sources of randomness. In this paper, we focus on the case where the noise directly affects the second component (the variable  $U$  of the FHN model), modelling the randomness of the conductance dynamics (see e.g. [3, 31]). This leads to a hypoelliptic stochastic differential equation (SDE), i.e., an SDE with a degenerate noise structure but a smooth transition density (the membrane potential component  $V$  is only indirectly impacted by noise). Measurements of the dynamics of a neuron, for example voltage-clamp experiments, provide observations of the membrane potential  $V$  at discrete times, while the recovery variable  $U$  is not observed. Such partially observed SDE is characterised by four parameters determining its behaviour, and it has been proved to be geometrically ergodic in [31].

Our goal is to estimate all four parameters of the stochastic FHN model in a Bayesian setting. To achieve this, we need to face and tackle several challenges. First, the solution and the transition density of the SDE are not available in an explicit closed-form, implying that numerical methods are needed to approximate them. Second, the SDE has a degenerate noise structure preventing using classical numerical schemes, such as the Euler-Maruyama discretisation, for likelihood approximation, as these schemes do not share this property [5, 9, 39, 47]. Third, the drift of the SDE is non-globally Lipschitz, and thus standard Itô-Taylor numerical methods do not converge to the true solution (in the strong mean-square sense) [23]. This problem may be overcome by using tamed [24, 51, 56, 57] or truncated [20, 22, 33, 34] variants. However, this type of methods still fails to be structure-preserving, meaning that they do not capture important properties of the model, such as its hypoellipticity (as discussed before), geometric ergodicity or oscillatory dynamics (see [5] and the references therein). Fourth, only partial observations of the system are available. This is difficult because the coordinate  $V$  alone is not Markovian, only the couple  $(V, U)$  shares this property.

In the frequentist framework, several estimation methods applicable to the stochastic FHN model have been proposed. They use adapted numerical schemes, and are based on estimation contrasts or optimal control theory. Among others, we can cite [25, 39, 47] in the case of complete observations, and [7, 9, 17, 26, 52] in the case of partial observations. In the Bayesian setting, we refer to [18, 19], where specific filtering algorithms, based on either the Taylor 1.5 scheme or the Euler-Maruyama scheme, are applied to derive the marginal posterior distributions of the FHN model parameters.

In this paper, we focus on the approximate Bayesian computation (ABC) method, see [35, 53] for a review. ABC is a simulation-based inference approach for Bayesian parameter estimation in complex mathematical models with unknown or intractable likelihood functions. In its basic acceptance-rejection version, it derives an approximate posterior distribution for the parameters of interest in three simple steps. First, a parameter value is sampled from a proposal distribution (e.g., the prior, for the simple acceptance-rejection ABC algorithm). Second, synthetic data are simulated from the model using the sampled parameter value. Third, such value is kept as a realisation from the approximate posterior distribution if the distance between some summary statistics of the reference data and the simulated data is smaller than some predefined tolerance level. This procedure is then repeated several times until, e.g., a certain number of sampled parameters are accepted. Among others, ABC has been proposed for inference in time series models [12, 27], state space models [36, 55] and SDE models, see e.g. [6, 10, 28, 30, 38, 43]. ABC has, however, not yet been applied to the stochastic FHN model, and, more generally, to hypoelliptic SDEs with locally Lipschitz drift coefficients, for which a convergent structure-preserving numerical scheme has been recently proposed in [5], as discussed below.

As ABC relies on running extensive simulations from the model under different parameter values (step two), the quality of its results depends strongly on the choice of the numerical scheme used for synthetic data generation. Due to its simple implementation and properties (e.g. Gaussian transition densities), the standard Euler-Maruyama scheme has been used for data-simulation from any SDE model (see, e.g., [43, 44, 45, 54]), including the stochastic FHN. However, in [6] it has been recently shown that such scheme may lead to drastic fails in the estimation, due to its non-structure-preserving nature. Instead, adopting a structure-preserving numerical scheme led to successful inference of the underlying parameters (from partially observed and hypoelliptic Hamiltonian-type SDEs, including oscillatory neuronal models). Following this, here we use a numerical splitting scheme for the data-simulation within ABC. This method was recently proposed in [5], and shown to be both *mean-square convergent* (of order 1) and *structure-preserving*, i.e., it retains the hypoellipticity, geometric ergodicity and oscillatory dynamics (amplitudes, frequencies and phases) of the stochastic FHN model, even for large time-discretisation steps. Hence, it is not only reliable, but also computationally efficient, which is particularly important when used within a simulation-based algorithm requiring millions of model simulations.

Besides the chosen numerical scheme, the performance of ABC also depends heavily on the choice of suitable summary statistics for the investigated data. When applying ABC to SDE models, it is particularly important to choose summary statistics that are sensitive to small parameter changes, and robust to the intrinsic stochasticity of the model. Inspired by [6], we propose to use statistics that are invariant for repeated simulations under the same parameter setting, namely the invariant density and spectral density, which are both based on the geometric ergodicity property of the stochastic FHN model. Moreover, we show how such choice leads to more accurate approximate posterior densities than those obtained using standard canonical summaries (e.g., mean, variance and auto-correlation).

Differently from [6], we do not embed here the structure-based summaries and structure-preserving numerical splitting scheme within a basic acceptance-rejection ABC algorithm. Indeed, such algorithm is very computationally wasteful, since parameter proposals are sampled from the prior distribution, which typically covers regions far away from the desired posterior distribution mass. Thus, many parameter candidates are rejected, after having spent computational time and effort to simulate the corresponding synthetic data, the summaries, and the distance measure. To improve this, several advanced ABC algorithms have been introduced over the years, such as sequential Monte Carlo (SMC) ABC (the gold standard state-of-art algorithm), Markov Chain Monte Carlo (MCMC) ABC, or sequential importance sampling (SIS) ABC, see [53] for a review. While such sequential ABC approaches have been applied to SDE models (see, e.g., [28, 43]), these works do typically not consider structure-based summaries and structure-preserving simulation methods. In this paper, we employ the structure-based ingredients within an efficient SMC-ABC scheme, yielding what we call *structure-based and preserving (SBP) SMC-ABC method*.

Our SBP SMC-ABC method succeeds in estimating all parameters of the stochastic FHN model from simulated data, even for a relatively small number of observations and a short observation time horizon. More interestingly, we also apply our method to real data consisting of membrane voltages recorded in an adult female Sprague Dawley rat [40]. We analyse four recordings, two obtained under a resting condition, and the other two under a stimulation setting. We show the coherence of the estimated posterior distributions with respect to these experimental setups. The fitted models are then validated through goodness-of-fit plots obtained by simulations under the estimated model.

While there exist only very few results fitting (or using) the FHN model to (for) real data applications, see e.g. [11] for a study on neural spike time data, and [32] for an investigation of vibration bearing data, this is, to the best of our knowledge, the first work on the considered stochastic FHN model investigating and fitting some real action potential data. This is a novel result which potentially broadens the models' scope and credibility. In addition, note that the proposed method is not limited to the stochastic FHN model, but it can be applied to any other ergodic SDE for which a computationally tractable structure-preserving numerical method is available.

The paper is organised as follows. In Section 2, we introduce the stochastic FHN model, and discuss its mathematical characteristics and structural properties. In Section 3, we introduce the SBP SMC-ABC method, detailing the algorithm, and all required key ingredients. In Section 4, we use this method to estimate the parameters of the stochastic FHN model from simulated data, under different observation settings. Section 5 details the results obtained on the real action potential data. Conclusions are reported in Section 6 and details on the algorithm are presented in the appendix.

## 2 Stochastic FHN model

The stochastic FHN model [5, 31] is a well-established prototype model of a single neuron (nerve cell), and dates back to the work of FitzHugh (1961) [16] and Nagumo (1962) [42]. It describes the generation of spikes (also called action potentials) at the intracellular level, and it can be obtained as solution to the following two-dimensional nonlinear SDE

$$d \underbrace{\begin{pmatrix} V(t) \\ U(t) \end{pmatrix}}_{:=X(t)} = \underbrace{\begin{pmatrix} \frac{1}{\epsilon} (V(t) - V^3(t) - U(t)) \\ \gamma V(t) - U(t) + \beta \end{pmatrix}}_{:=F(X(t))} dt + \underbrace{\begin{pmatrix} 0 \\ \sigma \end{pmatrix}}_{:=\Sigma} dW(t), \quad X(0) = X_0, \quad t \in [0, T], \quad (1)$$

where  $(W(t))_{t \in [0, T]}$  is a standard Wiener process, and the initial value  $X_0 = (V_0, U_0)^\top$  is a  $\mathbb{R}^2$ -valued random variable with bounded moments that is independent of  $(W(t))_{t \in [0, T]}$ . This is a common formulation of the model, where no noise enters directly into the  $V$ -component of SDE (1), and only the  $U$ -component is directly perturbed by it. The considered stochastic FHN model contains four parameters

$$\theta = (\epsilon, \gamma, \beta, \sigma), \quad (2)$$

where  $\beta, \gamma > 0$  describe the position and duration of an action potential, respectively,  $\epsilon > 0$  is a time scale separation parameter, and  $\sigma > 0$  defines the noise intensity. The  $V$ -component of system (1) models the evolution of the membrane voltage of the neuron, which can be measured with suitable recordings techniques (cf. Section 5), while the  $U$ -component is an unobserved recovery variable modelling the ion channel kinetics.

**Mathematical characteristics and structural properties** The drift  $F(X(t))$  of the FHN model (1) is not globally Lipschitz continuous. However, as it satisfies a *one-sided Lipschitz* condition [4, 5], the SDE has a unique strong (unknown) solution  $X(t) = (V(t), U(t))^\top$ ,  $t \in [0, T]$ , which is regular in the sense of [29], that is, it is defined on the whole time interval  $[0, T]$ , and paths of the process do not blow up to infinity in finite time. Moreover, such solution is a Markov process with (unknown) transition probability

$$P_t(\mathcal{A}, x) := \mathbb{P}(X(t) \in \mathcal{A} | X(0) = x), \quad \mathcal{A} \in \mathcal{B}(\mathbb{R}^2), \quad (3)$$

where  $\mathcal{B}(\mathbb{R}^2)$  denotes the Borel sigma-algebra on  $\mathbb{R}^2$ .

Furthermore, the process  $(X(t))_{t \in [0, T]}$  possesses two known structural properties. First, it is *hypoelliptic* [9, 31]. This means that the transition distribution, defined via (3), has a smooth density, even though the squared diffusion term

$$\Sigma \Sigma^\top = \begin{pmatrix} 0 & 0 \\ 0 & \sigma^2 \end{pmatrix}$$

is not of full rank, due to the degenerate noise structure of the FHN model (1). This property holds because the  $U$ -component (the model coordinate which is directly perturbed by the noise) enters into the first component of the drift  $F(X(t))$ .

Second, it is *geometrically ergodic* [1, 5, 31, 37], i.e., the distribution of the solution converges exponentially fast to a unique invariant distribution  $\eta$  for any initial value  $X_0$ , with  $\eta$  satisfying

$$\eta(\mathcal{A}) = \int_{\mathbb{R}^2} P_t(\mathcal{A}, x) \eta(dx), \quad \forall \mathcal{A} \in \mathcal{B}(\mathbb{R}^2), t \in [0, T].$$

In particular, if  $X_0$  has distribution  $\eta$ , then  $X(t)$  has distribution  $\eta$  for all  $t \in [0, T]$ . The existence of the invariant distribution  $\eta$  follows from the fact that there exists a Lyapunov function for system (1). For further details regarding the characteristics and structural properties of the stochastic FHN model, the reader is referred to [5, 31] and the references therein.

**Statistical setting** The stochastic FHN model (1) is only *partially observed* through discrete-time measurements of its first coordinate, the  $V$ -component. The measurements are made at equidistant times  $0 = t_0 < t_1 < \dots < t_{n-1} < t_n = T_{\text{obs}}$ , with  $t_i = i\Delta_{\text{obs}}$ ,  $\Delta_{\text{obs}} = T_{\text{obs}}/n$ ,  $i = 0, \dots, n$ , where  $T_{\text{obs}}$  and  $\Delta_{\text{obs}}$  denote the observation time horizon and time step, respectively. Here, the resulting observation vector is denoted by  $y = (V(t_i))_{i=0}^n$ . The goal of this paper is to derive an approximate posterior density (via the ABC method) for  $\theta$  (2), given the available observations  $y$ .

### 3 Structure-based and preserving SMC-ABC method

In this section, we introduce the proposed ABC methodology for parameter estimation of the stochastic FHN model (1). In Section 3.1, we recall the gold-standard SMC-ABC algorithm. In Section 3.2, we detail our proposed key ingredients of this algorithm, yielding the SBP SMC-ABC method. In Section 3.3, we provide implementation details.

#### 3.1 SMC-ABC algorithm

Let  $\pi(\theta)$  denote the prior density of an unknown parameter  $\theta$ ,  $p(y|\theta)$  the (unavailable) likelihood function of observations  $y$ , and  $\pi(\theta|y) \propto p(y|\theta)\pi(\theta)$  the desired posterior density. We now recall the basic acceptance-rejection ABC algorithm presented in the introduction, to introduce the relevant notation. Such scheme consists of three steps: (a) sample  $\theta'$  from the prior  $\pi(\theta)$ ; (b) conditioned on  $\theta'$ , simulate a synthetic dataset  $\tilde{y}_{\theta'}$  from the model; (c) keep the sampled value  $\theta'$  if the distance  $d(\cdot, \cdot)$  between the vector of summary statistics  $s(\cdot)$  of the observed and simulated data is smaller than a user-defined tolerance level  $\delta$ . This leads to an approximate posterior given by

$$\pi(\theta|y) \approx \pi_{\text{ABC}}^\delta(\theta|s(y)) \propto \int \mathbb{1}_{\{d(s(y), s(\tilde{y}_{\theta'})) < \delta\}} \pi(\theta) p(s(\tilde{y}_{\theta'})|\theta) ds.$$

As this ABC scheme is highly computationally inefficient, sequential ABC schemes have then been proposed to tackle this, with SMC-ABC being the most prominent one. The idea of SMC-ABC is to construct a sequence of  $r$  intermediate ABC posteriors using proposal samplers (also known as perturbation kernels), which are based on the  $N$  kept sampled values (called particles) at the previous iterations, and are thus continuously updated across iterations.

The first iteration of SMC-ABC coincides with acceptance-rejection ABC with a tolerance level  $\delta_1$ , after which all kept particles  $(\theta_1^{(1)}, \dots, \theta_1^{(N)})$  are given an equal weight  $w_1^{(j)} = 1/N$ ,  $j = 1, \dots, N$ , and a smaller threshold  $\delta_2$  is proposed. At each iteration  $r > 1$ , a value  $\theta_j$  is initially sampled from the set of kept particles  $(\theta_{r-1}^{(j)})_{j=1}^N$  at iteration  $r-1$  according to the corresponding weights  $(\omega_{r-1}^{(j)})_{j=1}^N$ , and then perturbed via a proposal sampler  $K_r(\cdot|\theta_j)$ , yielding  $\theta_j^*$ . Such value is then used to simulate new data  $\tilde{y}_{\theta_j^*}$ , and  $\theta_j^*$  is accepted if  $d(s(y), s(\tilde{y}_{\theta_j^*})) < \delta_r$ . After  $N$  particles are accepted, the unnormalised weights are updated as

$$\tilde{\omega}_r^{(j)} = \pi(\theta_r^{(j)}) / \sum_{l=1}^N \omega_{r-1}^{(l)} k_r(\theta_r^{(j)}|\theta_{r-1}^{(l)}), \quad j = 1, \dots, N,$$

and then normalised via  $\omega_r^{(j)} = \tilde{\omega}_r^{(j)} / \sum_{l=1}^N \tilde{\omega}_r^{(l)}$ . Here,  $k_r(a|b)$  denotes the density of the proposal sampler  $K_r(\cdot|b)$  evaluated in  $a$ .

The algorithm terminates when a stopping criterion is met, e.g.  $\delta_r$  becomes smaller than a desired final threshold  $\delta_{\text{final}}$ , a certain number of iterations is reached, or the acceptance rate of a particle becomes lower than 1.5% for two consecutive iterations [8]. Here, the algorithm terminates when a computational budget  $\text{Nsim}_{\text{max}} \gg 0$ , consisting of the desired total number of synthetic model simulations, is reached. The corresponding procedure is detailed in Algorithm 1. Note that, the computational budget  $\text{Nsim}_{\text{max}}$  is exceeded in the last iteration, so the used budget is larger than the specified one.

---

**Algorithm 1** SMC-ABC (with computational budget as stopping criterion)

**Input:** Observed dataset  $y$ , summary statistics  $s(\cdot)$ , prior distribution  $\pi(\theta)$ , perturbation kernel  $K_r$ , computational budget  $\text{Nsim}_{\text{max}}$  (i.e. maximum number of simulated datasets), number of kept samples in each iteration  $N$  and starting threshold  $\delta_1$

**Output:** Weighted samples from the SMC-ABC posterior distribution

---

```

1: Set  $r = 1$  and  $\text{Nsim} = 0$ .
2: for  $j = 1 : N$  do
3:   repeat
4:     Sample  $\theta_j$  from the prior  $\pi(\theta)$ .
5:     Conditioned on  $\theta_j$ , simulate a synthetic dataset  $\tilde{y}_{\theta_j}$  from the model,
   and set  $\text{Nsim} = \text{Nsim} + 1$ 
6:     Compute the summaries  $s(\tilde{y}_{\theta_j})$ .
7:     Calculate the distance  $D_j = d(s(y), s(\tilde{y}_{\theta_j}))$ .
8:   until  $D_j < \delta_1$ 
9:   Set  $\theta_1^{(j)} = \theta_j$ .
10:  Set  $\tilde{\omega}_1^{(j)} = 1, \omega_1^{(j)} = 1/N$ .
11: end for
12: while  $\text{Nsim} < \text{Nsim}_{\text{max}}$  do
13:   Set  $r = r + 1$ .
14:   Set  $\delta_r < \delta_{r-1}$ .
15:   for  $j = 1 : N$  do
16:     repeat
17:       Sample  $\theta_j$  from the weighted set  $\{(\theta_{r-1}^{(j)}, \omega_{r-1}^{(j)})_{j=1}^N\}$ .
18:       Perturb  $\theta_j$  to obtain  $\theta_j^*$  from  $K_r(\cdot|\theta_j)$ .
19:       Conditioned on  $\theta_j^*$ , simulate a synthetic dataset  $\tilde{y}_{\theta_j^*}$  from the model,
     and set  $\text{Nsim} = \text{Nsim} + 1$ 
20:       Compute the summaries  $s(\tilde{y}_{\theta_j^*})$ .
21:       Calculate the distance  $D_j = d(s(y), s(\tilde{y}_{\theta_j^*}))$ .
22:     until  $D_j < \delta_r$ 
23:     Set  $\theta_r^{(j)} = \theta_j^*$ .
24:     Set  $\tilde{w}_r^{(j)} = \pi(\theta_r^{(j)}) / \sum_{l=1}^N w_{r-1}^{(l)} k_r(\theta_r^{(j)}|\theta_{r-1}^{(l)})$ .
25:   end for
26:   Normalise the weights  $w_r^{(j)} = \tilde{w}_r^{(j)} / \sum_{l=1}^N \tilde{w}_r^{(l)}$ , for  $j = 1, \dots, N$ .
27: end while
28: Return a set of weighted parameters  $(\theta_r^{(1)}, \omega_r^{(1)}), \dots, (\theta_r^{(N)}, \omega_r^{(N)})$ .

```

---

## 3.2 Ingredients for the SMC-ABC algorithm

The SMC-ABC Algorithm 1 requires the choice of several key ingredients: the data simulation method, the summaries (along with a suitable distance measure), the prior distribution, the proposal sampler, the threshold levels, and the stopping criterion (here, the allowed maximum number of synthetic data simulations). Among all, the first two play a crucial role, as we would like to

apply an efficient numerical scheme capable of simulating from the model while preserving its key (structural) features, as well as to choose some summary statistics based on such properties. The chosen ingredients, detailed below, guarantee this, yielding the proposed SBP SMC-ABC method.

### 3.2.1 Structure-preserving synthetic data simulation with a splitting scheme

A crucial step within each simulation-based statistical method is the generation of synthetic datasets from the model for a given parameter value (cf. Lines 5 and 19 of Algorithm 1). As exact simulation schemes are not available for the stochastic FHN model (1), and standard Itô-Taylor type methods are non-convergent and/or non-preserving, we will simulate data from SDE (1) via the numerical splitting scheme recently proposed in [5].

This method is based on the idea of re-writing the drift coefficient  $F(X(t))$  of the stochastic FHN model (1) as the sum of a linear and a non-linear part, and then splitting the model into two exactly solvable subequations: a linear SDE, which shares the properties of hypoellipticity and geometric ergodicity, and a non-linear ordinary differential equation (ODE) capturing the non-globally Lipschitz polynomial term. The exact solutions of these two subequations are then derived and, in every iteration step, composed in a suitable way (here, according to the Strang approach, which considers “half-steps” of the ODE-solution). Due to the exact treatment of relevant and cleverly chosen subparts of the stochastic FHN model (1), the resulting (Strang) splitting scheme preserves its structural properties (hypoellipticity, geometric ergodicity, and oscillatory dynamics). The path generation procedure required in Algorithm 1 using this splitting scheme is briefly summarized in the following. For details regarding its construction and features, the interested reader is referred to [5].

For the model simulation, we consider an equidistant time grid  $0 = t_0 < t_1 < \dots < t_{n-1} < t_n = T$ , with  $t_i = i\Delta$ ,  $\Delta = T/n$ ,  $i = 0, \dots, n$ . We set  $\tilde{X}(t_0) := X_0$ , and denote by  $\tilde{X}(t_i)$  the splitting approximation of the solution  $X(t_i)$  of SDE (1) at time  $t_i$ . We then define the value

$$\kappa := 4\gamma/\epsilon - 1, \quad (4)$$

and the matrices involved in the exact solution of the linear SDE of the splitting framework, first the matrix exponential arising from the drift

$$E(t) := e^{-\frac{t}{2}} \begin{pmatrix} \cos(\frac{1}{2}\sqrt{\kappa}t) + \frac{1}{\sqrt{\kappa}} \sin(\frac{1}{2}\sqrt{\kappa}t) & -\frac{2}{\epsilon\sqrt{\kappa}} \sin(\frac{1}{2}\sqrt{\kappa}t) \\ \frac{2\gamma}{\sqrt{\kappa}} \sin(\frac{1}{2}\sqrt{\kappa}t) & \cos(\frac{1}{2}\sqrt{\kappa}t) - \frac{1}{\sqrt{\kappa}} \sin(\frac{1}{2}\sqrt{\kappa}t) \end{pmatrix}, \quad (5)$$

and second the corresponding covariance matrix

$$C(t) := \begin{pmatrix} c_{11}(t) & c_{12}(t) \\ c_{12}(t) & c_{22}(t) \end{pmatrix}, \quad (6)$$

where for  $t > 0$ ,

$$\begin{aligned} c_{11}(t) &= \frac{\sigma^2 e^{-t}}{2\epsilon\gamma\kappa} \left( -\frac{4\gamma}{\epsilon} + \kappa e^t + \cos(\sqrt{\kappa}t) - \sqrt{\kappa} \sin(\sqrt{\kappa}t) \right), \\ c_{12}(t) &= \frac{\sigma^2 e^{-t}}{\kappa\epsilon} \left( \cos(\sqrt{\kappa}t) - 1 \right), \\ c_{22}(t) &= \frac{\sigma^2 e^{-t}}{2\kappa} \left( \cos(\sqrt{\kappa}t) + \sqrt{\kappa} \sin(\sqrt{\kappa}t) - \frac{4\gamma}{\epsilon} + \kappa e^t \right). \end{aligned}$$

Finally, we define the function that solves exactly the ODE subequation of the splitting framework

$$h(x; t) := \begin{pmatrix} v/\sqrt{e^{-\frac{2t}{\epsilon}} + v^2 \left(1 - e^{-\frac{2t}{\epsilon}}\right)} \\ \beta t + u \end{pmatrix},$$

---

**Algorithm 2** Synthetic-data simulation of the FHN model (1) via Strang splitting [5]

**Input:** Initial value  $X_0$ , step-size  $\Delta$ , number of time discretisation steps  $n$ , model parameters  $\theta = (\epsilon, \gamma, \beta, \sigma)$  such that  $\kappa$  (4) is positive.

**Output:** Simulated path  $(\tilde{X}(t_i))_{i=0}^n$  of (1), with  $\tilde{y}_\theta = (\tilde{V}(t_i))_{i=0}^n$  used for inference

---

- 1: Compute the matrices  $E(\Delta)$  (5) and  $C(\Delta)$  (6).
  - 2: Set  $\tilde{X}(t_0) = X_0$ .
  - 3: **for**  $i = 1 : n$  **do**
  - 4: Generate  $\xi = (\xi_1, \xi_2)^\top$  from a bivariate normal distribution  $\mathcal{N}((0, 0)^\top, C(\Delta))$ .
  - 5: Set  $a = h(\tilde{X}(t_{i-1}); \Delta/2)$ .
  - 6: Set  $b = E(\Delta)a + \xi$ .
  - 7: Set  $\tilde{X}(t_i) = h(b; \Delta/2)$ .
  - 8: **end for**
  - 9: Return  $\tilde{X}(t_i)$ , for  $i = 0, \dots, n$  and use  $\tilde{y}_\theta = (\tilde{V}(t_i))_{i=0}^n$  for inference.
- 

for  $t > 0$ ,  $x = (v, u)^\top \in \mathbb{R}^2$ . A path of the FHN model (1) is then obtained via the Strang splitting scheme reported in Algorithm 2. Note that, this algorithm generates a trajectory for both coordinates  $\tilde{X}(t_i) = (\tilde{V}(t_i), \tilde{U}(t_i))^\top$ ,  $i = 0, \dots, n$ , but only the first coordinate  $\tilde{y}_\theta = (\tilde{V}(t_i))_{i=0}^n$  is used in the inferential procedure.

**Remark 1** *The requirement  $\kappa > 0$  in Algorithm 2 guarantees that the underlying stochastic harmonic oscillator (i.e. the embedded linear SDE) of the FHN model (1) entering into the splitting framework is weakly damped, yielding suitable oscillatory dynamics. In our experiments, the prior will be chosen to meet this condition, see Section 3.2.4. This requirement can be however relaxed, and we refer to [5] for the corresponding expressions for  $\kappa \leq 0$ .*

### 3.2.2 Structure-based data summaries

The summary statistics employed here (cf. Lines 6 and 20 of Algorithm 1) are the structure-based summaries proposed in [6] for geometrically ergodic SDEs. In particular, we consider the invariant density  $f_\theta$ , and the invariant spectral density  $S_\theta$  given by

$$S_\theta(\nu) = \int_{\mathbb{R}} R_\theta(\tau) \exp(-i2\pi\nu\tau) d\tau,$$

which is the Fourier transform of the auto-correlation function  $R_\theta(\tau)$  with respect to the frequency  $\nu$ . Here,  $f_\theta$  and  $S_\theta$  are estimated based on a given dataset  $\tilde{y}_\theta$ , via a standard kernel density estimator  $\hat{f}_{\tilde{y}_\theta}$  [48], and a smoothed periodogram estimator  $\hat{S}_{\tilde{y}_\theta}$  [49], respectively, yielding the summaries

$$s(\tilde{y}_\theta) := (\hat{f}_{\tilde{y}_\theta}, \hat{S}_{\tilde{y}_\theta}). \quad (7)$$

Thanks to the ergodicity of the SDE, and its preservation by the numerical splitting scheme in Algorithm 2, these summaries have several desired features. First, they can be well estimated from a single path  $\tilde{y}_\theta$  generated over a sufficiently long time horizon, instead of requiring a large number of repeated path simulations. Second, they are negligibly impacted by the (possibly unknown) initial value  $X_0$ . Moreover, as they are deterministic functions of the unknown parameter  $\theta$ , they fully characterise the asymptotic behaviour of the SDE (third), they are sensitive to small changes in  $\theta$  (fourth), and robust to repeated simulations under the same parameter values (fifth). These characteristics are illustrated in Figure 1, where the top panel reports two paths obtained under the same value of  $\theta$  (green and black trajectories), and a third path obtained under a slightly different value for  $\theta$  (red trajectory). The corresponding summaries, derived from these datasets, are shown in the bottom panels. While judging on the similarities/differences between the trajectories seems difficult, this becomes straightforward when looking at their summaries, which are indistinguishable only when estimated from data generated under the same parameter value.

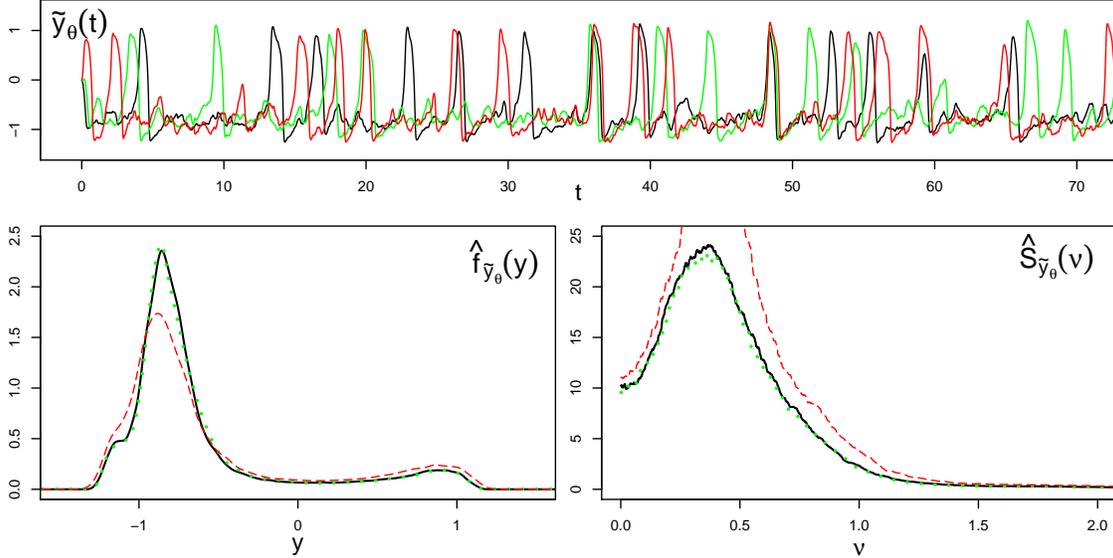


Figure 1: Top panel: Three paths of the observed component  $V(t)$  of SDE (1) generated with Algorithm 2. The parameter values are  $\theta = (0.1, 1.5, 0.8, 0.3)$  for the black and green paths and  $\theta = (0.1, 1.6, 0.8, 0.4)$  for the red path. Bottom panels: Corresponding estimated invariant density (left) and spectral density (right).

**Remark 2** A common alternative to the proposed summary statistics are what we call here canonical summaries, i.e., the mean, variance, skewness, kurtosis and auto-correlation up to a certain lag (here 5) of the original data and their differences, the latter to account for temporal dependence. This leads to a 18-dimensional summary vector which we denote by  $s_1(\tilde{y}_\theta)$  for a dataset  $\tilde{y}_\theta$ , and will only be considered in Section 4.2.

### 3.2.3 Choice of the distance measure

As a distance measure between the summaries of the observed dataset  $y$  and synthetic dataset  $\tilde{y}_\theta$  (cf. Lines 7 and 21 of Algorithm 1), we consider

$$d(s(y), s(\tilde{y}_\theta)) := \text{IAE}(\hat{S}_y, \hat{S}_{\tilde{y}_\theta}) + \alpha \text{IAE}(\hat{f}_y, \hat{f}_{\tilde{y}_\theta}), \quad (8)$$

where IAE is the *integrated absolute error* between two  $\mathbb{R}$ -valued functions  $g_1, g_2$  given by

$$\text{IAE}(g_1, g_2) := \int_{\mathbb{R}} |g_1(x) - g_2(x)| dx \geq 0. \quad (9)$$

The constant  $\alpha \geq 0$  in (8) corresponds to a *weight* guaranteeing that the two IAEs have a “comparable influence” on the distance  $d$  [6], which is relevant, as the invariant spectral density does not integrate to one. Here, the weight  $\alpha$  is chosen as the magnitude of the area below the estimated spectral density  $\hat{S}_y$  of the observed dataset  $y$ . Moreover, since  $\hat{S}_y, \hat{S}_{\tilde{y}_\theta}, \hat{f}_y$  and  $\hat{f}_{\tilde{y}_\theta}$  are estimated at discrete data points, the IAE (9) is approximated using rectangular integration.

**Remark 3** The Wasserstein distance between  $y$  and  $\tilde{y}_\theta$ , as proposed in [2], may be used instead of  $\text{IAE}(\hat{f}_y, \hat{f}_{\tilde{y}_\theta})$ , leading to marginal differences in the ABC posterior results (figures not shown). Moreover, as distance associated with the canonical summary statistics  $s_1(\cdot)$ , we consider the weighted Euclidean distance  $d_1(s_1(y), s_1(\tilde{y}_\theta)) = \|(s_1(y) - s_1(\tilde{y}_\theta)) / \alpha_1\|$ , where  $\alpha_1 \geq 0$  denotes a weight computed as mean absolute deviation of all summaries obtained during the first iteration of Algorithm 1 (i.e., the acceptance-rejection step).

### 3.2.4 Choice of prior distribution

The prior distribution is relevant in the acceptance-rejection step of Algorithm 1 ( $r = 1$ , see Line 4), and later on, when updating the weights at the  $r$ th iteration (cf. Line 24). Here, we use two sets of uniform priors

$$\epsilon \sim U(0.01, 0.5), \quad \gamma|\epsilon \sim U\left(\frac{\epsilon}{4}, 6\right) \quad \beta \sim U(0.01, 6), \quad \sigma \sim U(0.01, 1), \quad (10)$$

$$\epsilon \sim U(0.01, 1), \quad \gamma|\epsilon \sim U\left(\frac{\epsilon}{4}, 10\right), \quad \beta \sim U(0.01, 10), \quad \sigma \sim U(0.01, 3), \quad (11)$$

with the former used for the simulation study (in Section 4), and the latter for the real data study (in Section 5). Sampling  $\gamma$  from a uniform in  $(\epsilon/4, \cdot)$  guarantees that the condition  $\kappa > 0$  is met, see Remark 1.

### 3.2.5 Proposal Sampler

At the  $r$ th iteration,  $\theta_j$ , sampled from the weighted set  $\{(\theta_{r-1}^{(j)}, \omega_{r-1}^{(j)})_{j=1}^N\}$ , is perturbed via a proposal sampler  $K_r(\cdot|\theta_j)$  to obtain  $\theta_j^*$  (cf. Lines 17 and 18 of Algorithm 1.) Following [15], we choose  $K_r(\cdot|\theta_j)$  to be a multivariate normal perturbation kernel with mean  $\theta_j$  and covariance matrix given by twice the empirical weighted covariance matrix  $\hat{\Sigma}_{r-1}$  of the population  $\{(\theta_{r-1}^{(j)}, \omega_{r-1}^{(j)})_{j=1}^N\}$ , i.e.  $K_r(\cdot|\theta_j) = \mathcal{N}(\theta_j, 2\hat{\Sigma}_{r-1})$ . Hence,  $k_r(\cdot|\theta_{r-1}^{(l)})$  in Line 24 of Algorithm 1 is then the density of  $\mathcal{N}(\theta_{r-1}^{(l)}, 2\hat{\Sigma}_{r-1})$  evaluated in  $\theta_r^{(j)}$ , for  $j, l \in \{1, \dots, N\}$ . This Gaussian proposal sampler, which we call here **standard**, is obviously not the only possible choice. In Appendix A, we also consider the *optimal local covariance matrix* (**olcm**) Gaussian proposal sampler introduced in [15], and compare the impact of the two samplers on the estimation results.

### 3.2.6 Threshold levels

The first tolerance level  $\delta_1$ , needed in the first iteration of the SMC-ABC Algorithm 1 (the acceptance-rejection step), is obtained via an ABC pilot study as the  $p$ th percentile of  $10^4$  distances between the summaries of the reference data and of the datasets generated with parameter values sampled from the prior. All following tolerance levels  $\delta_r$ ,  $r > 1$ , are automatically chosen as the  $p$ th percentile of the  $N$  accepted distances computed at the previous iteration. Alternatively, one may choose the new tolerance levels as a percentile of all previous distances (both rejected and accepted), as done in [46], or via a procedure which, progressively, albeit slowly, reduces them while maintaining a reasonably high effective sample size (ESS), as suggested in [8].

## 3.3 Implementation details

All algorithms are implemented in the software environment R [50], using the package **Rcpp** [13], offering a seamless integration of R and C++. In particular, Algorithm 2 is coded in C++, which speeds up the synthetic data generation significantly, reducing thus the computational cost of the ABC procedure. Taking advantage of the independence between the  $N$  particles in the *for loops* in Algorithm 1, the code is parallelised using the packages **foreach** and **doParallel**. The R-package **mvnfast** [14] is used to draw samples from the multivariate normal perturbation kernel (using the function **rmvn**), and to compute its density (using the function **dmvn**).

The invariant spectral density  $S_{\tilde{y}_\theta}$  and invariant density  $f_{\tilde{y}_\theta}$  for a dataset  $\tilde{y}_\theta$ , used as summaries in (7), are estimated via the R-functions **spectrum** and **density**, respectively. Within the function **spectrum**, we avoid a logarithmic scale by setting **log**="no". Within the function **density**, we use the default value for the bandwidth **bw**, and set the number of points at which the density is estimated to **n**=  $10^3$ .

The number of particles accepted in each iteration is fixed to  $N = 10^3$  in all experiments. The threshold levels are chosen as the 50th percentile (median) of the simulated (for the pilot study) or accepted (for  $r \geq 1$ ) distances. All synthetic datasets (cf. Lines 5 and 19 of Algorithm 1) are generated with Algorithm 2, using the initial value  $X_0 = (0, 0)^\top$ , and the step-size  $\Delta = 0.02$ .

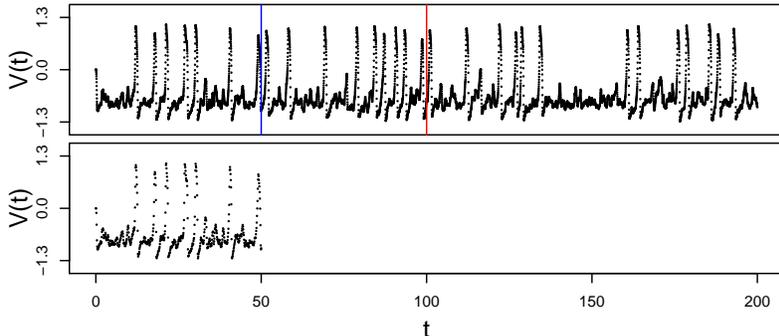


Figure 2: Top panel: Observed (simulated) dataset for  $\Delta_{\text{obs}} = 0.02$  and  $T_{\text{obs}} = 50$  (yielding  $n = 2500$  data points, up to the blue line),  $T_{\text{obs}} = 100$  ( $n = 5000$ , up to the red line), and  $T_{\text{obs}} = 200$  ( $n = 10^4$ ). Bottom panel: Observed (simulated) dataset for  $\Delta_{\text{obs}} = 0.08$  and  $T_{\text{obs}} = 50$  (yielding  $n = 625$  data points).

An easy-to-grasp implementation of the proposed SBP SMC-ABC method (i.e., Algorithm 1 with the ingredients presented in Section 3.2, including Algorithm 2 for synthetic data simulation) is available at [https://github.com/IreneTubikanec/SMC-ABC\\_FHN](https://github.com/IreneTubikanec/SMC-ABC_FHN). Moreover, a comprehensive R-package, providing the codes for running different sequential ABC schemes, with different choices of the summary statistics, distances and proposal samplers (see Section 4.2 and the appendix) will be available upon publication.

All case studies, based on both simulated and real data, have been run on multiple core High-Performance-Clusters at both the University of Klagenfurt, and the University of Warwick.

## 4 Inference of the FHN parameters from simulated data

In this section, we investigate the performance of the proposed statistical procedure for parameter estimation in the stochastic FHN model (1). In Section 4.1, we analyse the SBP SMC-ABC method, based on different simulated observed datasets. In Section 4.2, we study the impact of summary statistics on the ABC results, illustrating how the structure-based summaries may outperform the canonical ones. Moreover, in Appendix A, we also extensively investigate the impact of proposal samplers in terms of acceptance rate, ESS, quality of the inference and running times. Other ingredients in the procedure, such as the percentile in defining the threshold levels (in particular, 10, 25, 50%), and alternative distances between the summaries (namely, the Wasserstein distance between the invariant densities instead of the IAE) have shown to have a limited impact on the results, and are therefore not reported here.

Throughout this section (and the appendix), the underlying true parameter value of  $\theta$  (2) used to generate the respective observed dataset is the same as in [9], namely

$$\theta = (\epsilon, \gamma, \beta, \sigma) = (0.1, 1.5, 0.8, 0.3).$$

### 4.1 Inference under different observation settings

The performance of the proposed SBP SMC-ABC method is investigated under different observation regimes, depending on the choice of the observation time horizon  $T_{\text{obs}}$ , and the observation time step  $\Delta_{\text{obs}}$ . In particular, we consider  $T_{\text{obs}} = 50, 100, 200$ , and  $\Delta_{\text{obs}} = 0.02, 0.04, 0.08$ , yielding observed datasets with a size ranging from  $n = 10^4$  (for  $T_{\text{obs}} = 200$  and  $\Delta_{\text{obs}} = 0.02$ ) to only  $n = 625$  (for  $T_{\text{obs}} = 50$  and  $\Delta_{\text{obs}} = 0.08$ ) data points.

All observed datasets are obtained via Algorithm 2, using a simulation step-size  $\Delta = 10^{-4}$  and time horizon  $T = 200$ . They are then sub-sampled with the corresponding observation time step  $\Delta_{\text{obs}}$ , and cut to the respective observation length  $T_{\text{obs}}$ , as visualised in Figure 2. The

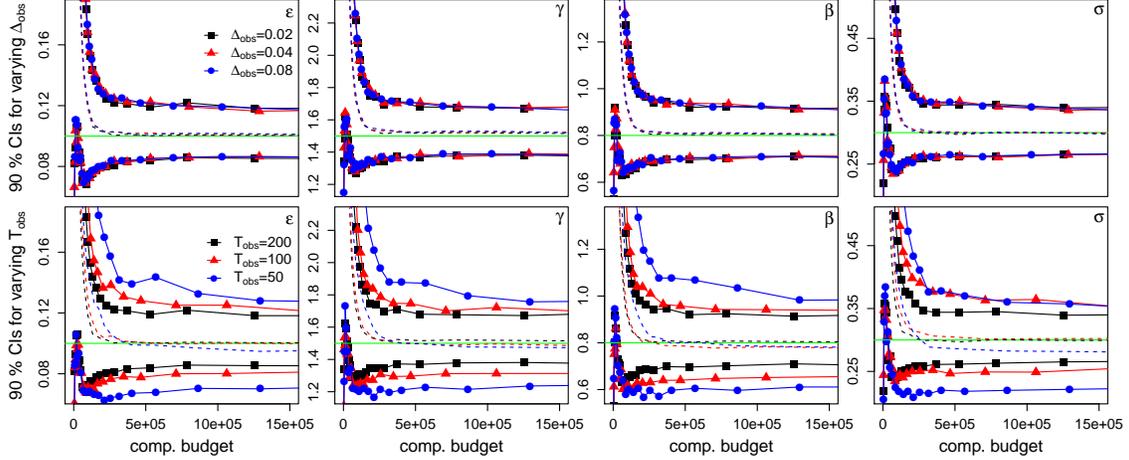


Figure 3: 90% CIs of the marginal posterior densities of the parameter vector  $\theta$  (2) of the stochastic FHN model (1) as a function of the computational budget. Top panels:  $T_{\text{obs}} = 200$ , and  $\Delta_{\text{obs}} = 0.02$  (black squares),  $0.04$  (red triangles),  $0.08$  (blue circles). Bottom panels:  $\Delta_{\text{obs}} = 0.02$ , and  $T_{\text{obs}} = 200$  (black squares),  $100$  (red triangles),  $50$  (blue circles). The green horizontal lines and the dashed lines indicate the true parameter values and the corresponding posterior medians, respectively.

synthetic datasets are instead obtained with  $T = T_{\text{obs}}$  and  $\Delta = 0.02$ , and then sub-sampled with the corresponding  $\Delta_{\text{obs}}$ . In this section, the computational budget of Algorithm 1 is set to  $N_{\text{sim}_{\text{max}}} = 10^6$  synthetic data simulations. Moreover, the results reported here are based on one trial per observation setting. However, we verified that they are consistent across repeated trials (figures not shown).

**ABC results** Our inferential results show that the observation time step  $\Delta_{\text{obs}}$  does not significantly influence the performance of the SBP SMC-ABC method, when choosing a (large enough) observation time horizon  $T_{\text{obs}}$ . This becomes evident when looking at the top panels of Figure 3, which show the 90% credible intervals (CIs) of the derived marginal ABC posterior densities of  $\theta$  (2) as a function of the computational budget, when fixing  $T_{\text{obs}} = 200$ , and considering  $\Delta_{\text{obs}} = 0.02$  (black squares),  $\Delta_{\text{obs}} = 0.04$  (red triangles) and  $\Delta_{\text{obs}} = 0.08$  (blue circles). This is also the case when choosing  $T_{\text{obs}} = 50$  or  $T_{\text{obs}} = 100$ .

In contrast to that, the observation time horizon  $T_{\text{obs}}$  does influence the inferential results when fixing the time-step  $\Delta_{\text{obs}}$ , with smaller  $T_{\text{obs}}$  yielding broader CIs. This can be observed in the bottom panels of Figure 3, where we fix  $\Delta_{\text{obs}} = 0.02$ , and consider  $T_{\text{obs}} = 200$  (black squares),  $T_{\text{obs}} = 100$  (red triangles), and  $T_{\text{obs}} = 50$  (blue circles). Similar considerations hold when choosing  $\Delta_{\text{obs}} = 0.04$ , or  $\Delta_{\text{obs}} = 0.08$ .

It is important to stress that the number of observed data points  $n$  is the same (for the same colours and shapes) in the top and bottom panels of Figure 3, i.e.,  $n = 2500$  (blue circles),  $n = 5000$  (red triangles) and  $n = 10000$  (black squares). However, the increased sample size leads to a more accurate ABC inference (in the sense of narrower CIs) only when increasing the time horizon. This is also confirmed by Table 1, where the ABC posterior means and standard deviations obtained under the different observation settings are reported. These findings are in agreement with the requirement  $T_{\text{obs}} \rightarrow \infty$  for contrast estimators of SDEs to be consistent and asymptotically normal (see, e.g., [9]).

Overall, all computed ABC posterior densities are unimodal, cover the true parameter values (green horizontal lines in Figure 3) with posterior medians (dashed lines in Figure 3) close to them (see also Table 1, where the respective posterior means are reported). Only for  $T_{\text{obs}} = 50$ , we observe that the ABC posterior medians (resp. the means) of the noise parameter  $\sigma$  are slightly off the true value of  $0.3$ .

Table 1: ABC posterior means and standard deviations of  $\theta$  (2), under different observation settings. The true parameter values are  $\epsilon = 0.1$ ,  $\gamma = 1.5$ ,  $\beta = 0.8$ , and  $\sigma = 0.3$ .

$T_{\text{obs}}$	$\Delta_{\text{obs}}$	$n$	ABC posterior means	ABC posterior stand. dev.s
200	0.02	$10^4$	(0.101, 1.523, 0.808, 0.300)	(0.010, 0.087, 0.062, 0.023)
200	0.04	5000	(0.101, 1.524, 0.808, 0.300)	(0.009, 0.086, 0.061, 0.022)
200	0.08	2500	(0.102, 1.531, 0.812, 0.301)	(0.010, 0.089, 0.064, 0.023)
100	0.02	5000	(0.101, 1.504, 0.787, 0.304)	(0.014, 0.119, 0.085, 0.034)
100	0.04	2500	(0.102, 1.508, 0.789, 0.305)	(0.014, 0.125, 0.089, 0.033)
100	0.08	1250	(0.101, 1.506, 0.787, 0.304)	(0.013, 0.116, 0.084, 0.033)
50	0.02	2500	(0.096, 1.480, 0.787, 0.284)	(0.017, 0.160, 0.113, 0.040)
50	0.04	1250	(0.098, 1.486, 0.789, 0.281)	(0.018, 0.166, 0.125, 0.039)
50	0.08	625	(0.098, 1.490, 0.793, 0.282)	(0.018, 0.171, 0.123, 0.041)

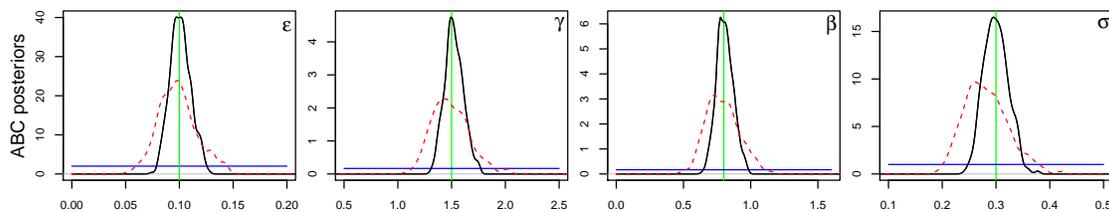


Figure 4: Marginal posterior densities of the parameter vector  $\theta$  (2) of the stochastic FHN model (1) (visualised over only a part of the prior domain). Black solid:  $T_{\text{obs}} = 200$ ,  $\Delta_{\text{obs}} = 0.02$ ,  $n = 10^4$ . Red dashed:  $T_{\text{obs}} = 50$ ,  $\Delta_{\text{obs}} = 0.08$ ,  $n = 625$ . The green vertical lines and the blue horizontal lines indicate the true parameter values and the prior densities (10), respectively.

In Figure 4, we report the marginal ABC posterior densities of  $\theta$  (2), derived from the two “uttermost” observation settings: (i)  $T_{\text{obs}} = 200$ ,  $\Delta_{\text{obs}} = 0.02$ ,  $n = 10^4$  (black solid lines); (ii)  $T_{\text{obs}} = 50$ ,  $\Delta_{\text{obs}} = 0.08$ ,  $n = 625$  (red dashed lines). In both scenarios, the marginal posterior densities show a clear update compared to the respective uniform prior densities (blue horizontal lines), and cover the true parameter values (green vertical lines), with the posteriors of scenario (i) being narrower than those of scenario (ii) (cf. the values for the corresponding ABC posterior standard deviations in Table 1).

## 4.2 Inference under different summary statistics

While the superiority and necessity of using structure-preserving numerical schemes over standard (non-preserving) simulation methods (such as Euler-Maruyama) within simulation-based algorithms have been intensively investigated in [6], no comparison between structure-based summaries (7) and “canonical” summaries (cf. Remark 2) has been provided yet. Here, we close this gap.

In Figure 5, we report the 90% CIs of the marginal posterior densities obtained via Algorithm 1 using the canonical (red lines) or the structure-based (black lines) summaries (the latter corresponding to the SBP SMC-ABC method). These CIs are plotted as a function of the computational budget, ranging up to 500000 (top panels), 1500000 (middle panels) or  $6 \cdot 10^6$  (bottom panels). For  $\beta$  and  $\sigma$ , the CIs obtained when using the canonical summaries are wider than those derived from the structure-based summaries up to circa 2-3 million model simulations, when they have approximately the same width. Moreover, using structure-based summaries yields CIs approaching the desired posterior region much faster, centered around the true parameter values for as few as 200-300 thousand simulations, and remaining almost unchanged when increasing the computational budget. This is not the case when using the canonical summaries. On the one hand,

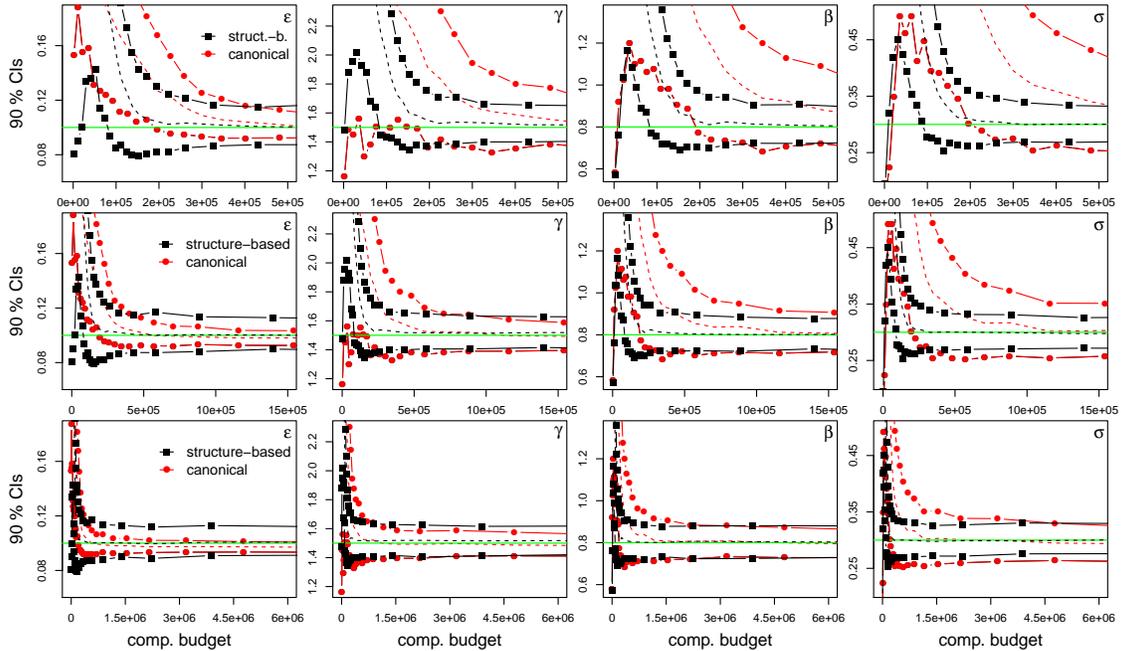


Figure 5: 90% CIs of the marginal posterior densities of the parameter vector  $\theta$  (2) of the stochastic FHN model (1) obtained from Algorithm 1 with either canonical (cf. Remark 2, red lines) or structure-based (cf. (7), black lines) summaries as a function of the computational budget. The green horizontal lines and the dashed lines indicate the true parameter values and the corresponding posterior medians, respectively.

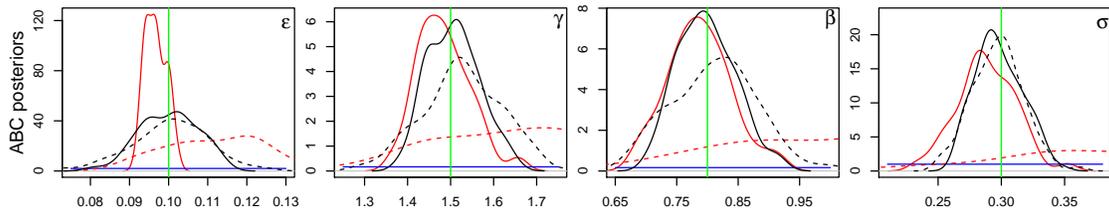


Figure 6: Marginal posterior densities of the parameter vector  $\theta$  (2) of the stochastic FHN model (1) obtained from the SMC-ABC Algorithm 1 with either canonical (cf. Remark 2, red lines) or structure-based (cf. (7), black lines) summaries when using approximately  $3 \cdot 10^5$  (dashed lines) and 11-13 million (solid lines) model simulations, respectively. The green vertical lines and the blue horizontal lines indicate the true parameter values and the prior densities, respectively.

they lead to posteriors with CIs for  $\epsilon$  and  $\sigma$  continuing shrinking when increasing the number of simulations (see the bottom panels). On the other hand, they yield posterior medians much larger than the true parameter values under small computational budgets (see the top panels), and below the true values under larger computational budgets (see in particular the bottom left panel for  $\epsilon$ ).

Similar observations can be made from Figure 6, where we plot the marginal posterior densities for a computational budget of approximately  $3 \cdot 10^5$  (dashed lines) and 12 million (solid lines) model simulations. While the structure-based summaries (black lines) yield satisfactory results even under a small computational budget, the canonical summaries (red lines) fail to do so. The results for  $\epsilon$  are particularly problematic, suggesting that the canonical summaries may be driving the posterior mass to a narrow and “wrong” area when increasing the computational budget, as also observed from the bottom left panel in Figure 5.

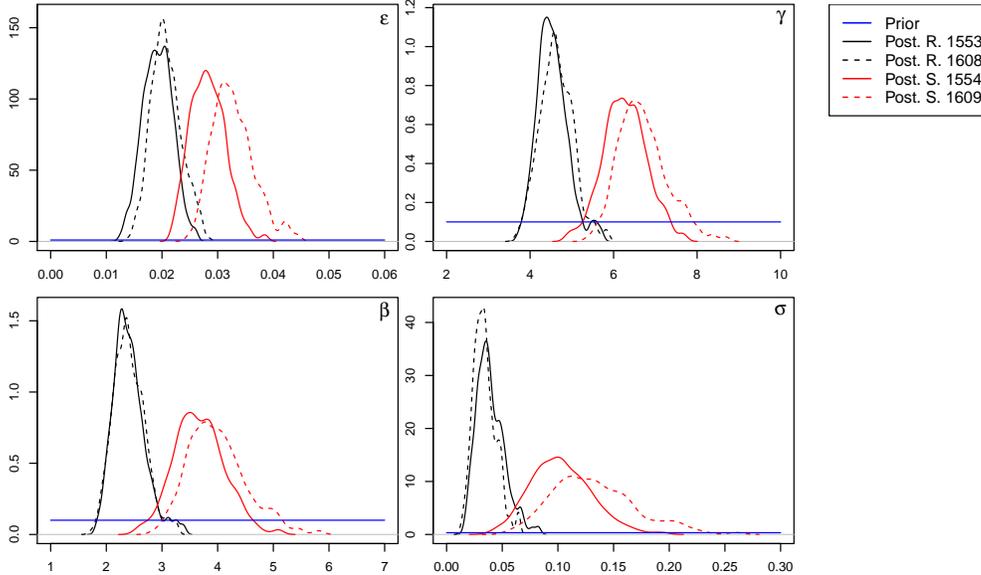


Figure 7: Marginal posterior densities of the parameter vector  $\theta$  (2) of the stochastic FHN model (1) fitted on the real action potential datasets *1553* (animal at rest, black solid lines), *1554* (animal stimulated, red solid lines), *1608* (animal at rest, black dashed lines), and *1609* (animal stimulated, red dashed lines), using the SBP SMC-ABC method. The horizontal blue lines indicate the uniform prior densities (11).

## 5 Inference of the FHN parameters from real action potential data

In this section, we infer the parameters  $\theta$  (2) of the stochastic FHN model (1) from real data consisting of membrane voltages recorded in rats.

### 5.1 Description of the data

The investigated data have been obtained via a neural recording from the 5th lumbar dorsal rootlet of an adult female Sprague Dawley rat, and are available at <https://data.mendeley.com/datasets/ybhwtngzmm/1>. A detailed description of the data and the corresponding experimental design are provided in [40]. In brief, these data consist of twenty recordings of length  $T_{\text{obs}} = 250$  ms with a sampling rate of 50 kHz (and thus  $\Delta_{\text{obs}} = 0.02$ ), ten obtained with the animal at rest (see the datasets in the folder “Resting”), and ten during stimulation (of the dermatome, activating the corresponding axons in the dorsal root, see the datasets in the folder “Cutaneous Stimulation”). Each recording consists of 5 channels of voltage data.

Here, we investigate four of the available twenty recordings, in particular the recordings *1553*, *1608* and *1554*, *1609*, stored in the folders “Resting” and “Cutaneous Stimulation”, respectively, the filenames corresponding to the times the recordings were started. In each recording, we focus on Channel 1, which represents the most distal electrode (i.e., the one closest to the tail). Similar results (not reported) are obtained when considering other channels and recordings.

### 5.2 Parameter inference

**ABC setup** The stochastic FHN model (1) is fitted to these four datasets using the SBP SMC-ABC method, with a computational budget of  $N_{\text{sim}_{\text{max}}} = 2 \cdot 10^6$ . All synthetic datasets are generated with step-size  $\Delta = \Delta_{\text{obs}} = 0.02$ , and time horizon  $T = T_{\text{obs}} = 250$ . Note that, the real datasets are centered close to zero, with a mean value of approximately 0.003 (the empirical means of the datasets *1553*, *1608*, *1554* and *1609* are 0.00277, 0.00265, 0.00271 and 0.00263, respectively). Therefore, in the following, all investigated datasets (real and synthetic) are centered.

Table 2: ABC posterior means, standard deviations and credible intervals of  $\theta$  (2), obtained under four different real action potential datasets.

case	ABC posterior means	ABC post. stand. dev.s	90% credible intervals
Rest. 1553	(0.019, 4.535, 2.407, 0.039)	(0.003, 0.368, 0.286, 0.013)	[0.015, 0.023], [3.991, 5.197], [2.007, 2.934], [0.022, 0.065]
Rest. 1608	(0.021, 4.629, 2.406, 0.034)	(0.003, 0.397, 0.281, 0.010)	[0.017, 0.026], [3.988, 5.313], [1.975, 2.871], [0.020, 0.051]
Stim. 1554	(0.028, 6.272, 3.670, 0.103)	(0.003, 0.521, 0.454, 0.028)	[0.023, 0.034], [5.425, 7.191], [2.964, 4.443], [0.062, 0.105]
Stim. 1609	(0.033, 6.701, 3.995, 0.133)	(0.004, 0.589, 0.529, 0.038)	[0.028, 0.039], [5.875, 7.758], [3.233, 5.003], [0.079, 0.204]

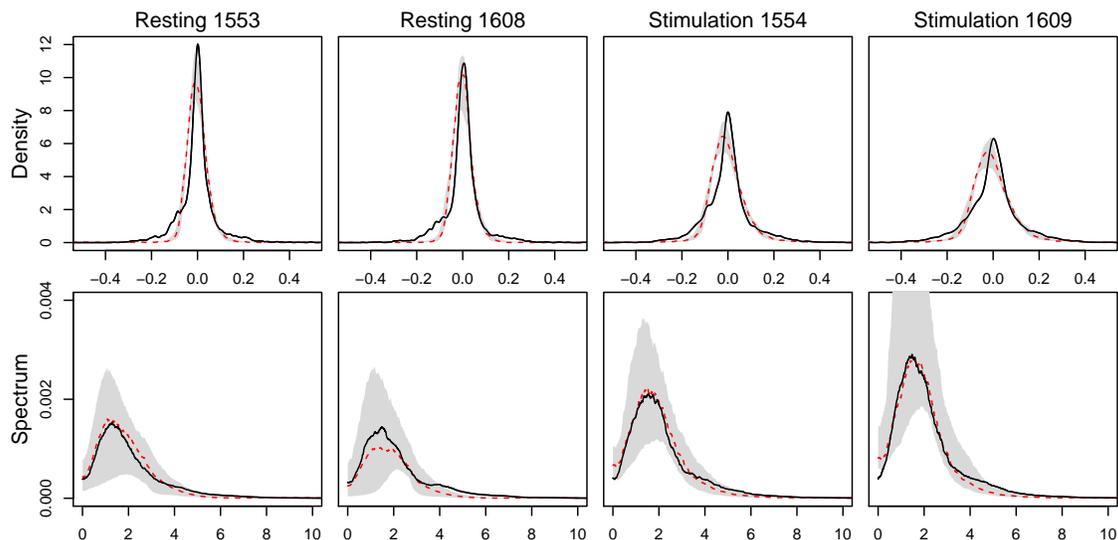


Figure 8: Structure-based summaries (7) of the action potential recordings 1553, 1554, 1608, 1609 (black solid lines) compared to summary domains (gray areas) obtained from synthetic datasets simulated under 50 kept ABC posterior samples (re-sampled using the corresponding weights). The red dashed lines represent summaries obtained under the corresponding ABC posterior means.

**Estimation results** Figure 7 shows the marginal ABC posterior densities obtained from the action potential recordings made when the animal was at rest (1553: black solid lines, 1608: black dashed lines), and when it was stimulated (1554: red solid lines, 1609: red dashed lines). The horizontal blue lines correspond to the underlying uniform prior densities, only a section of their domains being visualised. In addition, Table 2 reports relevant descriptive statistics of the obtained posterior distributions.

All marginal posterior densities are unimodal, notably different from the priors. Interestingly, the posteriors for the two “resting” (respectively “stimulating”) scenarios are similar, while they differ across these two cases, with the ABC marginal posterior densities of all parameters shifted towards larger values in the “stimulating” scenarios compared to the “resting” scenarios (see also Table 2). Moreover, their uncertainty also increases under the “stimulating” settings, as confirmed by the larger posterior standard deviations, and the wider 90% CIs reported in Table 2.

**Fitted FHN model** Fitting real data, we do not have ground truth parameters to compare our ABC results with. However, we check here whether the obtained results fit the model, which would advocate in favour of both the model and the inferential procedure. Figure 8 reports the estimated structure-based summary statistics (7) from the four investigated real action potential recordings (black solid lines) in comparison with the summaries estimated from synthetic datasets simulated under fifty parameter values sampled from the ABC posterior according to their weights (gray areas), and under the ABC posterior means reported in Table 2 (red dashed lines).

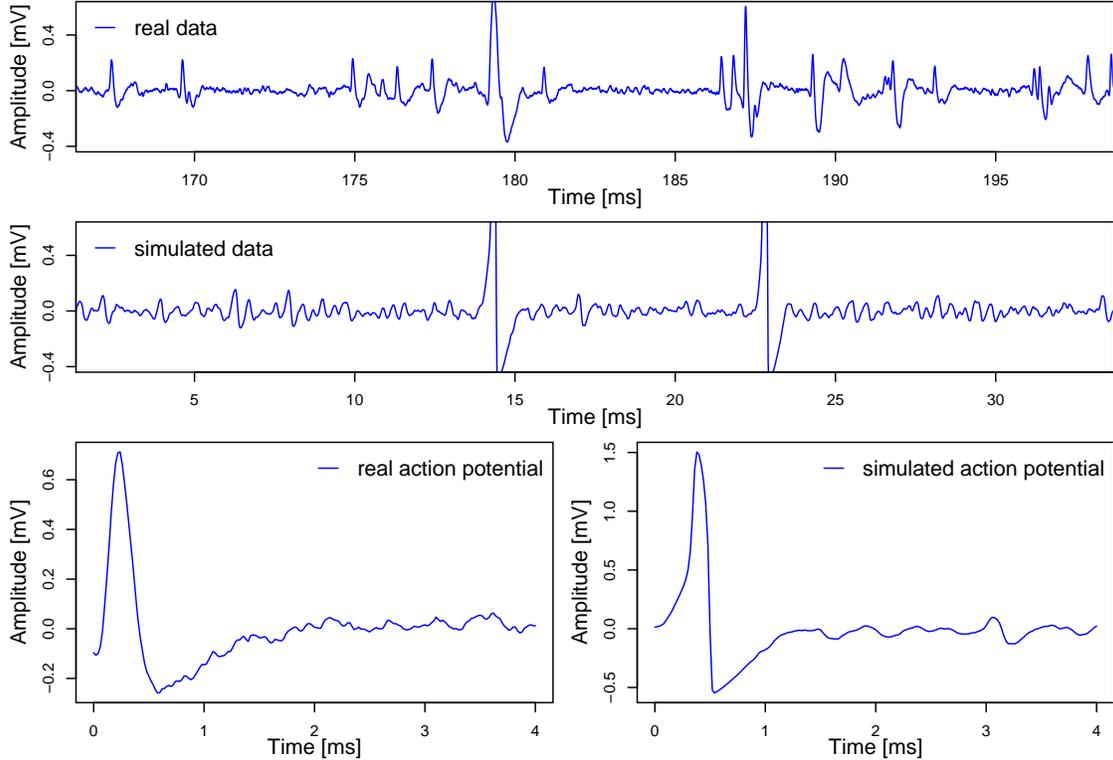


Figure 9: Exemplary 30 ms data snippet from recording *1553* (top panel) compared to a simulated dataset obtained under the corresponding ABC posterior means (middle panel). Exemplary real action potential (bottom left panel) taken from recording *1553* compared to a simulated action potential (bottom right panel) taken from the synthetic dataset visualised in the middle panel.

We observe a good fit of the summary statistics from real and synthetic data in all considered scenarios, with the gray bands for the densities being narrower than those for the spectral densities. Note also that the peaks of the real data spectral densities (resp. densities) for the “stimulating” scenarios are higher (resp. lower) than those for the “resting” scenarios, both features being captured by the summaries coming from the fitted model.

This nice fit is also observed in Figure 9. There, we compare both an exemplary 30 ms snippet from the real action potential recording *1553* (top panel) with a 30 ms simulated dataset generated under the corresponding posterior means reported in Table 2 (middle panel), and an action potential taken from the real data recording *1553* (bottom left panel) with another taken from the reported simulated dataset (bottom right panel). The synthetic action potential closely resembles the shape of the real action potential, including the depolarization (rapid increase of voltage), repolarization (rapid decrease of voltage), and refractory (time during which the cell cannot repeat an action potential) periods.

## 6 Conclusion

We propose a SMC variant of the SBP ABC method initially presented in [6] for a different class of SDEs, with the goal of estimating the parameters of the stochastic FHN model. This is a well known SDE used to describe the spiking activity of a single neuron at the intracellular level. The proposed SBP SMC-ABC method relies on summary statistics based on the ergodicity property of the model [31], and on a recently proposed structure-preserving numerical splitting scheme for reliable and efficient synthetic data generation [5].

While existing estimation methods often require complete observations, a non-degenerate noise structure, globally Lipschitz coefficients or a very complicated architecture, our method overcomes the associated difficulties, is easy-to-grasp, and straight-forward to implement and to parallelise. Moreover, in contrast to other inference procedures (which, e.g. require to fix the time scale separation parameter  $\epsilon$ , see [9]), the SMC SBP-ABC algorithm succeeds in simultaneously estimating all four model parameters of the stochastic FHN model, even for fewer observations and relatively small observation time horizons.

While there exist very few real data studies using the FHN model, the presented analysis of real action potential is the first for the considered stochastic version of the model. Here, we investigate real data recorded under different experimental conditions (resting and cutaneous stimulation) in an adult female rat [40]. The inferred posterior distributions are in agreement with the different experimental settings, and lead to satisfactory goodness-of-fit plots, suggesting the goodness of the considered model and the proposed estimation approach.

Besides successfully deriving marginal posterior densities from both synthetic and real data, we also discuss the influence of different key ingredients of ABC. For example, the proposed structure-based summaries avoid possible biases which may be instead introduced by the canonical ones, while other choices of percentiles (e.g. the 25th) or distances (e.g. the Wasserstein) have a limited impact on the results. In terms of proposal samplers, which one to choose depends on the available computational budget. Indeed, `olcm` (and other guided sequential ABC schemes proposed in [46]) leads to more accurate posterior inference than `standard` (the perturbation kernel considered here) for low/moderate computational budgets, with the differences becoming negligible as the computational budget increases. On the other hand though, `olcm` and the guided sequential schemes are characterised by more involved proposal means and covariances, introducing a computational overhead with respect to `standard`.

Finally, we note that the proposed method can be applied to any other ergodic SDE for which a computationally efficient and reliable numerical method for synthetic data generation is available. The exploration of more complex stochastic neuronal models is certainly of interest, and will be considered in future works.

## References

- [1] M. Ableidinger, E. Buckwar, and H. Hinterleitner. A stochastic version of the Jansen and Rit neural mass model: analysis and numerics. *J. Math. Neurosci.*, 7(8), 2017.
- [2] E. Bernton, P.J. Jacob, M. Gerber, and C.P. Robert. Approximate Bayesian Computation with the Wasserstein Distance. *J. Royal Stat. Soc. B*, 81(2):235–269, 2019.
- [3] S. Bonaccorsi and E. Mastrogiacomio. Analysis of the stochastic FitzHugh-Nagumo system. *Infim. Dimens. Anal. Quantum Probab. Relat. Top.*, 11(03):427–446, 2008.
- [4] C.-E Bréhier, D. Cohen, and G. Giordano. Splitting schemes for FitzHugh–Nagumo stochastic partial differential equations. *Discrete Cont Dyn-B*, 29(1):214–244, 2024.
- [5] E. Buckwar, A. Samson, M. Tamborrino, and I. Tubikanec. A splitting method for SDEs with locally Lipschitz drift: Illustration on the FitzHugh-Nagumo model. *Appl. Numer. Math.*, 179:191–220, 2022.
- [6] E. Buckwar, M. Tamborrino, and I. Tubikanec. Spectral density-based and measure-preserving ABC for partially observed diffusion processes. An illustration on Hamiltonian SDEs. *Stat. Comput.*, 30(3):627–648, 2020.
- [7] Q. Clairon and A. Samson. Optimal control for parameter estimation in partially observed hypoelliptic stochastic differential equations. *Comput. Stat.*, 37:2471–2491, 2022.
- [8] P. Del Moral, A. Doucet, and A. Jasra. An adaptive sequential Monte Carlo method for approximate Bayesian computation. *Stat. Comput.*, 22:1009–1020, 2012.

- [9] S. Ditlevsen and A. Samson. Hypoelliptic diffusions: filtering and inference from complete and partial observations. *J. Royal Stat. Soc. B*, 81(2):361–384, 2019.
- [10] S. Ditlevsen, M. Tamborrino, and I. Tubikanec. Network inference in a stochastic multi-population neural mass model via approximate Bayesian computation. *arXiv*, 2306.15787v2, 2023.
- [11] R. O. Doruk and L. Abosharb. Estimating the Parameters of Fitzhugh-Nagumo Neurons from Neural Spiking Data. *Brain Sci.*, 9:364, 2019.
- [12] C. Drovandi, A. Pettitt, and R. McCutchan. Exact and approximate Bayesian inference for low integer-valued time series models with intractable likelihoods. *Bayesian Anal.*, 11:325–352, 2016.
- [13] D. Eddelbuettel and R. François. Rcpp: seamless R and C++ integration. *J. Stat. Soft.*, 40(8):1–18, 2011.
- [14] M. Fasiolo. *An introduction to mvnfast. R package version 0.2.8*. University of Bristol, 2014.
- [15] S. Filippi, P. B. Barnes, J. Cornebise, and M. P. H. Stumpf. On optimality of kernels for approximate Bayesian computation using sequential Monte Carlo. *Stat. Appl. Genet. Mol. Biol.*, 12(1):87–107, 2013.
- [16] R. FitzHugh. Impulses and physiological states in theoretical models of nerve membrane. *Biophys. J.*, 1(6):445–466, 1961.
- [17] A. Gloter and N. Yoshida. Adaptive estimation for degenerate diffusion processes. *Elec. J. Stat.*, 15:1424–1472, 2021.
- [18] M.M. Graham, A.H. Thiery, and A. Beskos. Manifold markov chain monte carlo methods for bayesian inference in diffusion models. *J. Roy. Stat. Soc. B*, 84:1229–1256, 2022.
- [19] S. Grazi, F. van der Meulen, M. Mider, and M. Schauer. Bayesian inference for SDE models: a case study for an excitable stochastic-dynamical model. *NextJournal*, 2020.
- [20] Q. Guo, W. Liu, X. Mao, and R. Yue. The partially truncated Euler–Maruyama method and its stability and boundedness. *Appl. Numer. Math.*, 115:235–251, 2017.
- [21] 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(4):500–544, 1952.
- [22] M. Hutzenthaler and A. Jentzen. Numerical approximations of stochastic differential equations with non-globally Lipschitz continuous coefficients. *Mem. Am. Math. Soc.*, 236, 2012.
- [23] M. Hutzenthaler, A. Jentzen, and P. E. Kloeden. Strong and weak divergence in finite time of Euler’s method for stochastic differential equations with non-globally Lipschitz continuous coefficients. *Proc. R. Soc. A*, 467:1563–1576, 2011.
- [24] M. Hutzenthaler, A. Jentzen, and P. E. Kloeden. Strong convergence of an explicit numerical method for SDEs with nonglobally Lipschitz continuous coefficients. *Ann. Appl. Probab.*, 22(4):1611–1641, 2012.
- [25] Y. Iguchi and A. Beskos. Parameter inference for hypo-elliptic diffusions under a weak design condition. *submitted*, 2023.
- [26] Y. Iguchi, A. Beskos, and M. Graham. Parameter estimation with increased precision for elliptic and hypo-elliptic diffusions. *submitted*, 2022.
- [27] A. Jasra. Approximate Bayesian computation for a class of time series models. *Int Stat. Rev.*, 83:405–435, 2015.

- [28] P. Jovanovski, A. Golightly, and U. Picchini. Towards data-conditional simulation for abc inference in stochastic differential equations. *arXiv*, preprint, 2023.
- [29] R. Khasminskii and F. C. Klebaner. Long term behavior of solutions of the Lotka-Volterra system under small random perturbations. *Ann. Appl. Probab.*, 11(3):952–963, 2001.
- [30] T. Kypraios, P. Neal, and D. Prangle. A tutorial introduction to Bayesian inference for stochastic epidemic models using approximate Bayesian computation. *Math Biosci.*, 287:42–53, 2017.
- [31] J.R. León and A. Samson. Hypocoelliptic stochastic FitzHugh-Nagumo neuronal model: mixing, up-crossing and estimation of the spike rate. *Ann. Appl. Probab.*, 28(4):2243–2274, 2018.
- [32] C. López, W. Zhong, S. Lu, F. Cong, and I. Cortese. Stochastic resonance in an underdamped system with FitzHugh-Nagumo potential for weak signal detection. *Journal of Sound and Vibration*, 411:34–46, 2017.
- [33] X. Mao. The truncated Euler–Maruyama method for stochastic differential equations. *J. Comput. Appl. Math.*, 290:370–384, 2015.
- [34] X. Mao. Convergence rates of the truncated Euler–Maruyama method for stochastic differential equations. *J. Comput. Appl. Math.*, 296:362–375, 2016.
- [35] J.M. Marin, P. Pudlo, C.P. Robert, and R. Ryder. Approximate Bayesian computational methods. *Stat Comput*, 22:1167–1180, 2012.
- [36] G.M. Martin, McCabern B.P.M., D.T. Frazier, W. Maneesoonthorn, and C.P. Robert. Auxiliary likelihood-based approximate Bayesian computation in state space models. *J Comput. Graph. Stat.*, 28:1–31, 2019.
- [37] J. C. Mattingly, A. M. Stuart, and D. J. Higham. Ergodicity for SDEs and approximations: locally Lipschitz vector fields and degenerate noise. *Stoch. Process. Their Appl.*, 101(2):185–232, 2002.
- [38] P. Mayband, I. Bojak, and R. Everitt. Fast approximate Bayesian inference for stable differential equation models. *submitted*, 2017.
- [39] A. Melnykova. Parametric inference for hypoelliptic ergodic diffusions with full observations. *Stat. Inf. Stoc. Proc.*, 23:595–635, 2020.
- [40] B. Metcalfe, A. Hunter, J. Graham-Harper-Cater, and J. Taylor. A dataset of action potentials recorded from the L5 dorsal rootlet of rat using a multiple electrode array. *Data in Brief*, 33:106561, 2020.
- [41] C. Morris and H. Lecar. Voltage oscillations in the barnacle giant muscle fiber. *Biophys. J.*, 35:193–213, 1981.
- [42] J. Nagumo, S. Arimoto, and S. Yoshizawa. An active pulse transmission line simulating nerve axon. *Proc. IRE*, 50(10):2061–2070, 1962.
- [43] U. Picchini. Inference for SDE models via Approximate Bayesian Computation. *J Comput. Graph. Stat.*, 23:1080–1100, 2014.
- [44] U. Picchini and Forman J. L. Accelerating inference for diffusions observed with measurement error and large sample sizes using approximate Bayesian computation. *J. Stat. Comput. Simul.*, 86(1):195–213, 2016.
- [45] U. Picchini and A. Samson. Coupling stochastic EM and approximate Bayesian computation for parameter inference in state-space models. *Comput. Stat.*, 33(1):179–212, 2018.

- [46] U. Picchini and M. Tamborrino. Guided sequential ABC schemes for intractable Bayesian models. *Accepted for publication in Bayesian Analysis. Preprint at ArXiv*, 2206.12235, 2022.
- [47] Y. Pokern, A. M. Stuart, and P. Wiberg. Parameter estimation for partially observed hypoelliptic diffusions. *J. Royal Stat. Soc. B*, 71(1):49–73, 2009.
- [48] O. Pons. *Functional Estimation for Density, Regression Models and Processes*. World Scientific Publishing, Singapore, 2011.
- [49] B. Quinn, I. Clarkson, and R. McKilliam. On the periodogram estimators of periods from interleaved sparse, noisy timing data. *IEEE Stat. Signal Processing Workshop*, pages 232–235, 2014.
- [50] R Development Core Team. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria, 2011.
- [51] S. Sabanis. Euler approximations with varying coefficients: the case of superlinearly growing diffusion coefficients. *Ann. Appl. Probab.*, 26(4):2083–2105, 2016.
- [52] A. Samson and M. Thieullen. A contrast estimator for completely or partially observed hypoelliptic diffusion. *Stoc. Proc. Appl.*, 122:2521–2552, 2012.
- [53] S. A. Sisson, Y. Fan, and M. Beaumont. *Handbook of Approximate Bayesian Computation*. Chapman & Hall/CRC Handbooks of Modern Statistical Methods. CRC Press, Taylor & Francis Group, 2018.
- [54] L. Sun, C. Lee, and J. A. Hoeting. Parameter inference and model selection in deterministic and stochastic dynamical models via approximate Bayesian computation: Modeling a wildlife epidemic. *Environmetrics*, 26(7):451–462, 2015.
- [55] A. Tancredi. Approximate Bayesian inference for discretely observed continuous-time multi-state models. *Biometrics*, 75:966–977, 2019.
- [56] M. V. Tretyakov and Z. Zhang. A fundamental mean-square convergence theorem for SDEs with locally Lipschitz coefficients and its applications. *SIAM J. Numer. Anal.*, 51(6):3135–3162, 2013.
- [57] Z. Zhang and H. Ma. Order-preserving strong schemes for SDEs with locally Lipschitz coefficients. *Appl. Numer. Math.*, 112:1–16, 2017.

## A Impact of the proposal kernel

In this section, we first define the Gaussian `olcm` proposal sampler introduced in [15], and then compare its performance with the Gaussian `standard` proposal sampler used throughout the paper (and also proposed in [15]). Alternative sequential algorithms, such as sequential important sampling ABC (SIS-ABC), and several guided sequential ABC schemes (both SIS-ABC and SMC-ABC), recently proposed in [46], have also been considered. As their performance is comparable to that of the SMC-ABC algorithm with the `olcm` sampler (and better when using canonical summaries, which however we discourage to choose), we do not report them here.

### A.1 Gaussian optimal local covariance matrix (`olcm`) sampler

The Gaussian `standard` proposal sampler  $K_r(\cdot|\theta_j) = \mathcal{N}(\theta_j, 2\hat{\Sigma}_{r-1})$  considered throughout the paper is characterised by a mean  $\theta_j$  sampled from the weighted set  $\{(\theta_{r-1}^{(j)}, \omega_{r-1}^{(j)})_{j=1}^N\}$ , and a covariance matrix which is not specific for  $\theta_j$ , as it is computed by using all kept particles at the previous iteration (see Section 3.2.5). The Gaussian `olcm` proposal sampler instead, is a perturbation kernel with a covariance matrix specific for the sampled  $\theta_j$ . Among the  $N$  accepted

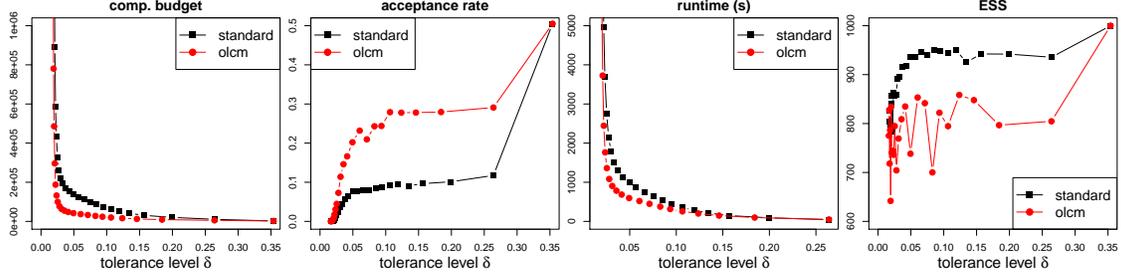


Figure 10: Computational budget, acceptance rate, runtime (in seconds), and ESS as a function of the tolerance level  $\delta$  when using the SBP SMC-ABC algorithm with either the **standard** (black lines) or the **olcm** (red lines) sampler.

particles at iteration  $r - 1$ , consider the subset of  $N_0 \leq N$  weighted particles whose corresponding distances (between the simulated and the reference data) are also smaller than  $\delta_r$ , i.e.

$$\{\tilde{\theta}_{r-1,l}, \mu_{r-1,l}\}_{1 \leq l \leq N_0} = \left\{ \left( \theta_{r-1}^{(j)}, \frac{w_{r-1}^{(i)}}{\mu_{r-1}} \right), \text{ such that } d(s(y), s(y_{\theta_{r-1}^{(j)}})) < \delta_r, \quad j = 1, \dots, N \right\},$$

where  $\bar{\mu}_{r-1}$  is a normalisation constant such that  $\sum_{l=1}^{N_0} \mu_{r-1,l} = 1$ . As we choose the new threshold level  $\delta_r$  as the  $p$ th percentile of the  $N$  accepted distances at iteration  $r - 1$  (see Section 3.2.6), we already know that the set above will contain the corresponding  $N_0 = p\%N$  particles. Then, at iteration  $r$ , the **olcm** proposal is given by  $K_r^{\text{olcm}}(\cdot|\theta_j) = \mathcal{N}(\theta_j, \Sigma_{\theta_j}^{\text{olcm}})$ , where  $\Sigma_{\theta_j}^{\text{olcm}} = \sum_{l=1}^{N_0} \mu_{r-1,l} (\tilde{\theta}_{r-1,l} - \theta_j)(\tilde{\theta}_{r-1,l} - \theta_j)^\top$ . Accepted particles are given unnormalized weights  $\tilde{w}_r^{(j)}$  as in Line 24 of Algorithm 1, with  $k_r(\cdot|\theta_{r-1}^{(l)})$  being now the density of  $\mathcal{N}(\theta_{r-1}^{(l)}, \Sigma_{\theta_{r-1}^{(l)}}^{\text{olcm}})$ . As the constructed covariance matrix  $\Sigma_{\theta_j}^{\text{olcm}}$  is now specific for each particle  $\theta_j$ , it will have to be computed for every  $j$ th particle,  $j = 1, \dots, N$ , introducing thus a small but non-negligible overhead in the computations.

## A.2 Inference under different proposal samplers

In this section, we investigate the impact of the Gaussian **standard** and **olcm** proposal samplers on the inference results, considering not only the obtained marginal ABC posterior densities, but also other quantities of interest. In particular, in Figure 10, we report the computational budget, acceptance rates (estimated as  $N$  divided by the number of sampled particles at iteration  $r$ ), the ESS (approximated at iteration  $r$  as  $1/(\sum_{j=1}^N (w_r^{(j)})^2)$ ), and the algorithm runtimes as a function of the tolerance level  $\delta$ . Remember that the first iteration of the algorithm consists on acceptance-rejection ABC, so the perturbation kernel does not play a role there. Moreover, as the threshold levels  $\delta$  are automatically decreased across iterations (with  $\delta_1 \approx 0.35$ ), the algorithm is “evolving” (in terms of successive iterations) towards smaller values of  $\delta$ , so from right to left in the plots of Figure 10. By looking at the first panel, we can see that **olcm** needs fewer model simulations than **standard** to achieve a certain tolerance level, which is consistent with the higher acceptance rates (second panel) across all algorithm iterations, as soon as the proposal samplers are used, i.e. from the second iteration. For example, **olcm** and **standard** achieve an approximate threshold level of 0.025 and 0.075, respectively, for a computational budget of  $10^5$  model simulations, requiring approximately 296000 and 888000 model simulations to achieve a threshold around 0.02. This is also consistent with the lower runtimes (obtained when running the algorithms on a desktop machine Intel Core i7-9700 CPU 3GHz RAM 32 GB) of **olcm** than **standard** (third panel). When looking at the ESS (fourth panel), a number in  $[1, N]$  measuring the effectiveness of the sampler, i.e., how many particles are “relevant” at each iteration, so the higher the better, we can see that, while high for both perturbation kernels, **standard** has higher ESS than **olcm** throughout the entire run of the algorithm.

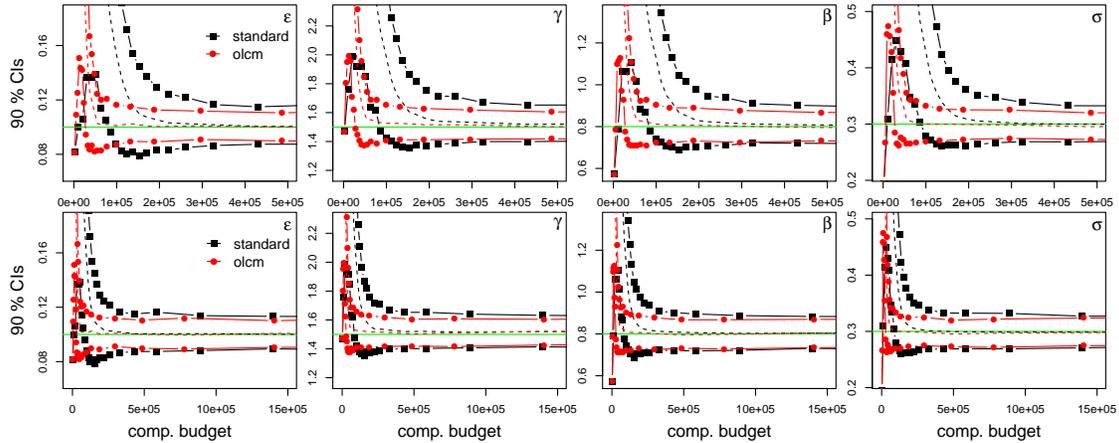


Figure 11: 90% CIs of the marginal posterior densities of  $\theta$  (2) of the stochastic FHN model (1) obtained from the SBP SMC-ABC algorithm with the **standard** (black lines) and **olcm** (red lines) samplers as a function of the computational budget. The green horizontal lines and the dashed lines indicate the true parameter values and the corresponding posterior medians, respectively.

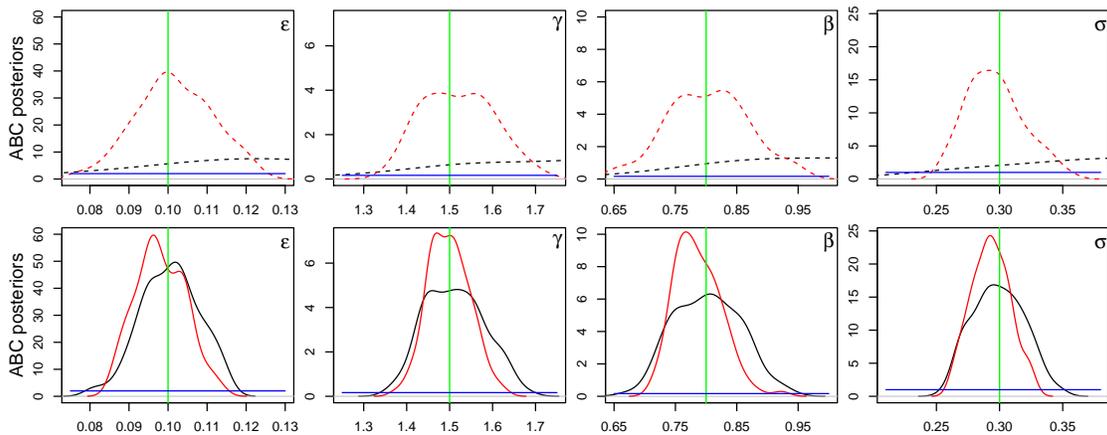


Figure 12: Marginal posterior densities of the parameter vector  $\theta$  (2) of the stochastic FHN model (1) obtained from the SBP SMC-ABC algorithm with **standard** (black lines) and **olcm** (red lines) when using a computational budget of approximately  $10^5$  (dashed lines, top panels) and 9-13 million (solid lines, bottom panels) model simulations, respectively. The green vertical lines and the blue horizontal lines indicate the true parameter values and the prior densities, respectively.

While these quantities are important and of interest, what matters the most is obviously the quality of the inference. In Figure 11, we report the 90% CIs (between the solid lines) and the median (dashed lines) of the marginal posterior densities obtained using the **standard** (black lines) and the **olcm** (red lines) proposal samplers, as a function of the computational budget. The choice of the sampler influences the inferential results, especially for a low/moderate computational budget. Indeed, **olcm** requires many fewer model simulations than **standard** to obtain narrow CIs which are centered around the true parameters, see the top panels in Figure 11. When the computational budget increases, such differences, while still present, become negligible. This can also be observed in Figure 12, where we report the SBP SMC-ABC marginal posteriors (black lines for **standard** and red lines for **olcm**) obtained with a computational budget of  $10^5$  (top panels, dashed lines) and 9-13 millions (bottom panels, solid lines) model simulations.