



Autonomous learning of nonlocal stochastic neuron dynamics

Tyler E. Maltba¹ · Hongli Zhao² · Daniel M. Tartakovsky³

Received: 22 November 2020 / Revised: 6 September 2021 / Accepted: 2 October 2021
© The Author(s), under exclusive licence to Springer Nature B.V. 2021

Abstract

Neuronal dynamics is driven by externally imposed or internally generated random excitations/noise, and is often described by systems of random or stochastic ordinary differential equations. Such systems admit a distribution of solutions, which is (partially) characterized by the single-time joint probability density function (PDF) of system states. It can be used to calculate such information-theoretic quantities as the mutual information between the stochastic stimulus and various internal states of the neuron (e.g., membrane potential), as well as various spiking statistics. When random excitations are modeled as Gaussian white noise, the joint PDF of neuron states satisfies exactly a Fokker-Planck equation. However, most biologically plausible noise sources are correlated (colored). In this case, the resulting PDF equations require a closure approximation. We propose two methods for closing such equations: a modified nonlocal large-eddy-diffusivity closure and a data-driven closure relying on sparse regression to learn relevant features. The closures are tested for the stochastic non-spiking leaky integrate-and-fire and FitzHugh-Nagumo (FHN) neurons driven by sine-Wiener noise. Mutual information and total correlation between the random stimulus and the internal states of the neuron are calculated for the FHN neuron.

Keywords Stochastic neuron model · Colored noise · Method of distributions · Nonlocal · Equation learning

Introduction

Computational neuroscience often employs the theory of nonlinear dynamical systems to quantitatively understand complex neuronal processes (Izhikevich 2007). A key aspect of the dynamical systems approach is the ability to discover relevant properties of brain function without

accounting for every detail. It comprises coupled ordinary differential equations (ODEs), with the number of ODEs determining the dimensionality of a dynamical system. Examples of such neuron representations are the one-dimensional leaky integrate-and-fire (LIF) model (Gerstner and Kistler 2002), two-dimensional FitzHugh-Nagumo (FHN) (FitzHugh 1961) and Morris-Lecar Morris and Lecar (1981) models, and four-dimensional Hodgkin-Huxley model (Hodgkin and Huxley 1952). To account for apparent randomness of neuron behavior (e.g., due to synaptic, internal, and/or channel noise), such models include a stochastic noise/source term. To account for parametric uncertainty, the coefficients in such models might be treated as random variables. In both cases, such models admit distributions of solutions that are (partially) characterized by single-time joint probability density functions (PDFs) of system states. These PDFs can be either used directly to characterize the ensemble behavior of large populations of neurons (Gerstner et al. 2014) or post-processed to compute interspike-interval (ISI) distributions (Greenwood and Ward 2016; Iolov et al. 2014), first exit time (FET) distributions (Alzubaidi and Shardlow 2014; Tuckwell and Wan 2005), and the mutual

Tyler E. Maltba: Supported by National Science Foundation Graduate Research Fellowship under Grant No. DGE 1752814.

Hongli Zhao: Supported by UC Berkeley SURF L&S Program.

Daniel M. Tartakovsky: Supported in part by Air Force Office of Scientific Research under award number FA9550-21-1-0381, and by a gift from TOTAL.

✉ Daniel M. Tartakovsky
tartakovsky@stanford.edu

¹ Department of Statistics, UC Berkeley, Berkeley, CA 94720, USA

² Department of Mathematics, UC Berkeley, Berkeley, CA 94720, USA

³ Department of Energy Resources Engineering, Stanford University, Stanford, CA 94305, USA

information between different system states, or between a system state and a stochastic input (Rieke 1997).

Synaptic noise reflects randomness in the release of neurotransmitters by synapses and random inputs from other neurons. Internal noise originates from the operations of ionic channels, and channel noise stems from the random switching of such channels.¹ Regardless of its genesis, a random source in ODEs (e.g., ionic current in the LIF and FHN models) is frequently treated as Gaussian white noise (Greenwood and Ward 2016). While mathematically and computationally convenient, Gaussian white noise is biologically implausible since it is unbounded and uncorrelated. For that reason, sine-Wiener (SW) channel noise—a bounded and correlated (colored) non-Gaussian process—has been used to model signal transmission (Kang et al. 2013), resonance (Liang et al. 2010), and detection (Liang et al. 2010) in nervous systems; to induce spike death in excitable systems (Guo et al. 2012); and to detect a weak periodic signal in the FHN neuron (Yao and Ma 2018). While our experiments focus solely on SW noise, the general methodology presented in “[Method of Distributions](#)” and “[Numerics](#)” Sections is equally applicable to arbitrary colored noise processes that have Lebesgue measurable and almost surely bounded paths. Another common choice, among many others, is the Ornstein-Uhlenbeck process, as seen in Kang et al. (2020); Guo (2011).

One of the attractive features of Gaussian white noise is that the resulting ODE-based neuron model can be viewed not only as a Langevin system, but also as a formal stochastic differential equation (SDE) whose solution has a joint PDF that satisfies exactly the Fokker-Planck equation (Risken and Frank 1996; Øksendal 2003). This deterministic advection-diffusion partial differential equation (PDE), whose dimensionality is determined by the number of state variables/ODEs in the underlying stochastic model, is no longer exact in the presence of colored noise. A joint PDF of neuron states—or other quantities such as ISI or FET—can be estimated via (multilevel) Monte Carlo simulations (Haskell et al. 2001; Rosenbaum 2001; Alzubaidi and Shardlow 2014; Giles 2008), but this procedure requires significant computational resources and sheds little light on the probabilistic dynamics of the neuron. More efficient methods of stochastic computation, e.g., polynomial chaos expansions and stochastic finite elements, do not provide a computational speed-up in this setting because colored noise is characterized by a short correlation length, leading to the so-called curse of dimensionality (see, e.g., the references in Wang et al. (2013)). The

¹ Due to the large number of these channels, channel noise is often ignored by arguing that fluctuations average out (Yamakou et al. 2019). However, channel noise in and of itself can change the behavior of neurons (White et al. 2000; Zhou et al. 2020), hence, it is important to study its effects on a neuron’s dynamics.

method of distributions (Tartakovsky and Gremaud 2015) can be used to derive a PDF equation for the stochastic solution of an ODE driven by colored noise (Wang et al. 2013; Maltba et al. 2018). In the limit of the vanishing correlation length, i.e., when colored noise reduces to white noise, such a PDF equation becomes exact, reducing to the corresponding Fokker-Planck equation. Otherwise, it requires a closure approximation to be computable.

We propose two alternative approaches to construct such closures: theory-based nonlocal large-eddy-diffusivity (LED) closures (Kraichnan 1987; Wang et al. 2013; Maltba et al. 2018) and a data-driven closure relying on statistical inference, i.e., dictionary learning of relevant derivative features in PDF equations (Bakarji and Tartakovsky 2021). To the best of our knowledge, both the method of distributions and statistical learning of PDF equations from a few Monte Carlo runs have not been used in computational neuroscience before. We investigate the relative performance of these strategies on two stochastic neuron models with varying degrees of complexity, the stochastic non-spiking leaky integrate-and-fire (NS-LIF) and FHN models. While standard LIF models are by far the simplest biologically plausible neuron models, which are advantageous for simulating large coupled networks, they are limited in their ability to capture the basic topology of a neuron’s excitability. The FHN model overcomes this limitation and captures key features of the Hodgkin-Huxley model. Our stochastic analogues of the NS-LIF and FHN models account for SW channel noise.

In “[Method of distributions](#)” Section, we present the method of distributions in the context of Langevin-type systems driven by colored noise, and describe alternative approaches to the construction of closures for resulting PDF equations. Our numerical strategies for solving these equations are discussed in “[Numerics](#)” Section. The method’s performance and errors associated with the proposed closures are investigated for the stochastic NS-LIF and FHN models in “[Numerical experiments](#)” Section. This section also contains our strategy for computing mutual information between channel noise (or stochastic input current) and a neuron’s membrane potential. Main conclusions derived from this study are summarized in “[Conclusion](#)” Section.

Method of distributions

Consider a system of random ordinary differential equations (RODEs),

$$\frac{d \mathbf{x}(t, \omega)}{d t} = \mathbf{v}(\mathbf{x}(t, \omega), t; \xi(t, \omega)), \quad \mathbf{x}(0, \omega) = \mathbf{x}^0(\omega), \quad (1)$$

that is to be solved on a time interval $(0, T_f]$ and holds for almost every $\omega \in \Omega$ in an appropriate probability space $(\Omega, \mathcal{F}, \mathbb{P})$. The solution $\mathbf{x}(t, \omega) = [x_1(t, \omega), \dots, x_N(t, \omega)]^\top : (0, T_f] \times \Omega \rightarrow \mathbb{R}^N$ is an \mathbb{R}^N -valued stochastic process, and the initial state \mathbf{x}^0 is an N -dimensional random vector with the joint probability density function (PDF) $f_{\mathbf{x}^0}(\mathbf{X}_0) : \mathbb{R}^N \rightarrow \mathbb{R}^+$. The given deterministic function $\mathbf{v} = [v_1, \dots, v_N]^\top : \mathbb{R}^N \rightarrow \mathbb{R}^N$, parameterized with a set of N_{par} random coefficients $\xi(t, \omega) = [\xi_1(t, \omega), \dots, \xi_{N_{\text{par}}}(t, \omega)]^\top$, satisfies conditions guaranteeing the existence of a unique path-wise solution $\mathbf{x}(t, \omega)$ (see for instance Strand 1970; Neckel and Rupp 2013; Asai and Kloeden 2016). The zero-mean random processes $\xi(t, \omega)$ are characterized by a prescribed single-point joint PDF $f_\xi(\Xi; t)$ and a two-point covariance matrix \mathbf{C}_ξ with elements $C_{ij} \equiv \mathbb{E}[\xi_i(t_1, \omega)\xi_j(t_2, \omega)]$, where $i, j = 1, \dots, N_{\text{par}}$ and $t_1, t_2 \in (0, T_f]$. Here and below we use $\mathbb{E}[\cdot]$ and $\langle \cdot \rangle$ interchangeably to denote the ensemble mean.

Let $\mathbf{X} = [X_1, \dots, X_N]^\top \in \mathbb{R}^N$ be a variable in the system's phase space. At any given time t , the state of the system is (partially) characterized by the joint probability $F_{\mathbf{x}}(\mathbf{X}; t) \triangleq \mathbb{P}[\mathbf{x}(t, \omega) \leq \mathbf{X}]$. If this single-point joint cumulative distribution function $F_{\mathbf{x}}(\mathbf{X}; t)$ is differentiable with respect to all X_i , then the system is described by the corresponding single-point joint PDF $f_{\mathbf{x}}(\mathbf{X}; t)$. Our goal is to derive a deterministic equation for the PDF $f_{\mathbf{x}}(\mathbf{X}; t)$ of the random solution $\mathbf{x}(t, \omega)$ to (1).

The RODE solution $\mathbf{x}(t, \omega)$ has paths that are continuously differentiable with derivatives that are at most Hölder continuous. Hence, paths of the driving sample processes (e.g., ξ) of (1) are at most Hölder continuous and need only be Lebesgue measurable and almost surely bounded so that (1) can be interpreted in the sense of Carathéodory. Therefore, the random coefficients $\xi_i(t, \omega)$ are correlated in time, i.e., $C_{ij} \sim \delta(t_1 - t_2)$ where $\delta(\cdot)$ is the Dirac delta function. If the coefficients $\xi_i(t, \omega)$ are uncorrelated in time (i.e., white noise), then (1) should be rewritten as a formal SDE (Øksendal 2003) for which $f_{\mathbf{x}}(\mathbf{X}; t)$ satisfies exactly a Fokker-Planck (advection-diffusion) equation.

Systems of RODEs can always be rewritten as a formal multidimensional SDE by treating the coefficients $\xi_i(t, \omega)$ as states and adding their governing SDEs to the system of equations (Neckel and Rupp 2013). Doing so would give a (possibly much) higher-dimensional Fokker-Planck equation for the joint density of $[x_1, \dots, x_N, \xi_1, \dots, \xi_{N_{\text{par}}}]$, significantly increasing computational cost when compared to the method of distributions, which seeks only a governing equation for $f_{\mathbf{x}}(\mathbf{X}; t)$. Such an equation is henceforth referred to as a PDF equation; it is not unique and its derivation requires a closure approximation (e.g., Wang et al. 2013 and the references therein).

In general, correlated inputs render PDF equations nonlocal in both phase space and time [ibid]. We focus on two interconnected challenges: numerical treatment of nonlocal PDF equations and statistical inference (i.e., learning) of the relevant closure terms in such equations from a few Monte Carlo realizations of (1).

We use the implementation of the method of distribution that treats random noise $\xi(t)$ exactly. This is in contrast to the implementations based on truncated expansions of random inputs, e.g., in a Karhunen-Loève series (Venturi et al. 2012, 2013). Such expansions are often not appropriate for colored noise $\xi(t)$ due to the presence of short correlation lengths. Derivation of a PDF equation starts by defining an auxiliary functional (Tartakovsky and Gremaud 2015)

$$\Pi(\mathbf{X}, t; \omega) \triangleq \delta(\mathbf{x}(t, \omega) - \mathbf{X}) = \prod_{i=1}^N \delta(x_i(t, \omega) - X_i). \tag{2}$$

At any time t , its ensemble mean over all realizations of \mathbf{x} is the PDF $f_{\mathbf{x}}(\mathbf{X}; t)$:

$$\mathbb{E}[\Pi] \triangleq \int_{\mathbb{R}^N} \delta(\mathbf{Y} - \mathbf{X}) f_{\mathbf{x}}(\mathbf{Y}; t) \, d\mathbf{Y} = f_{\mathbf{x}}(\mathbf{X}; t).$$

We show in “Appendix 1” that $\Pi(\mathbf{X}, t; \omega)$ obeys, in the sense of distributions, the linear conservation law

$$\frac{\partial \Pi}{\partial t} + \nabla_{\mathbf{X}} \cdot [\mathbf{v}(\mathbf{X}, t; \omega) \Pi] = 0, \quad \Pi(\mathbf{X}, 0) = \delta(\mathbf{x}^0(\omega) - \mathbf{X}). \tag{3}$$

This equation describes, e.g., advection, in the phase space \mathbf{X} , of a passive scalar (with concentration Π) in the random velocity field \mathbf{v} . In that context, its stochastic averaging/homogenization has been the subject of multiple investigations (Kraichnan 1987).

A typical homogenization procedure employs the Reynolds decomposition to represent the random inputs and outputs as the sums of their means and zero-mean fluctuations around these means. For example, $\mathbf{v} = \langle \mathbf{v} \rangle + \mathbf{v}'$ where

$$\langle \mathbf{v} \rangle(\mathbf{X}, t) \triangleq \int_{\mathbb{R}^{N_{\text{par}}}} \mathbf{v}(\mathbf{X}, t; \Xi) f_\xi(\Xi) \, d\Xi \quad \text{and} \quad \langle \mathbf{v}' \rangle = \mathbf{0}.$$

Likewise, the Reynolds decomposition of Π is $\Pi = f_{\mathbf{x}} + \Pi'$. Substituting these decompositions into (3) and taking the ensemble mean yields an unclosed PDF equation

$$\frac{\partial f_{\mathbf{x}}}{\partial t} + \nabla_{\mathbf{X}} \cdot (\langle \mathbf{v} \rangle f_{\mathbf{x}}) + \nabla_{\mathbf{X}} \cdot \langle \mathbf{v}' \Pi' \rangle = 0, \quad f_{\mathbf{x}}(\mathbf{X}; 0) = f_{\mathbf{x}^0}(\mathbf{X}), \tag{4}$$

with vanishing free-space boundary conditions (though others may be considered). This equation is unclosed

because it contains the unknown flux term $\langle \mathbf{v}'\Pi' \rangle$. Following (Maltba et al. 2018; Barajas-Solano and Tartakovsky 2016), we consider three alternative versions of the LED closure: local, semi-local, and nonlocal. In all three formulations, the terminal problem (“Appendix 1.2”)

$$\frac{d\boldsymbol{\chi}}{ds} = \langle \mathbf{v}(\boldsymbol{\chi}, s) \rangle, \quad \boldsymbol{\chi}(t) = \mathbf{X}, \quad (5)$$

and its associated flow $\boldsymbol{\chi}(s) \equiv \boldsymbol{\Phi}(s; \mathbf{X}, t)$ play an essential role. Whenever possible, we omit ω for notational convenience.

Local LED closure

The local closure is constructed by a direct application of the classical LED approach in fluid dynamics (Kraichnan 1987). Under that approach, the cross covariance $\langle \mathbf{v}'\Pi' \rangle(\mathbf{X}, t)$ is split into advection and diffusion components,

$$\langle \mathbf{v}'\Pi' \rangle \approx \mathcal{V}(\mathbf{X}, t) f_{\mathbf{x}}(\mathbf{X}; t) - \mathcal{D}(\mathbf{X}, t) \nabla_{\mathbf{X}} f_{\mathbf{x}}(\mathbf{X}; t), \quad (6)$$

where the LED drift velocity \mathcal{V} and the LED diffusion tensor \mathcal{D} are given in (56) and (57), respectively. Underpinning this closure is the assumption that the density $f_{\mathbf{x}}$ and its spatial derivatives vary slowly over a correlation-length interval $(t - \tau, t)$. Therefore, this closure is expected to be valid for systems with short correlation lengths.

Nonlocal LED closure

The nonlocal closure does not rely on the short-correlation assumption to the extent of the local closure and is second-order accurate in $\sigma\tau$, i.e., in the product of the noise intensity and correlation length (see van Kampen 1976 and “Appendix 1.2”). Under the nonlocal closure, the cross-covariance term is approximated by

$$\langle \mathbf{v}'\Pi' \rangle \approx - \int_0^t \mathcal{J}(s; \mathbf{X}, t) \nabla_{\Phi} \cdot \left(\langle \mathbf{v}'(\mathbf{X}, t) \mathbf{v}'^{\top}(\boldsymbol{\Phi}(s; \mathbf{X}, t), s) \rangle f_{\mathbf{x}}(\boldsymbol{\Phi}(s; \mathbf{X}, t); s) \right) ds, \quad (7a)$$

where

$$\mathcal{J}(s; \mathbf{X}, t) = \exp\left(- \int_s^t \nabla_{\boldsymbol{\chi}} \cdot \langle \mathbf{v}(\boldsymbol{\chi}(r), r) \rangle dr\right). \quad (7b)$$

For a given dimension N , the computational complexity of the nonlocal PDF method is dominated by the work of solving the terminal problem (5) and evaluating the Jacobian determinant (7b). The latter is dramatically simplified if the flow divergence $\nabla_{\boldsymbol{\chi}} \cdot \langle \mathbf{v} \rangle$ is linear in $\boldsymbol{\chi}$, which occurs in such classical test problems as the nonlinear pendulum and the Duffing oscillator; analyses of such problems, e.g.,

Cho et al. (2013), are likely to give an overly optimistic assessment of PDF methods. The computational experiments reported in “Numerical experiments” Section are chosen to have nonlinear flow divergence $\nabla_{\boldsymbol{\chi}} \cdot \langle \mathbf{v} \rangle$.

Semi-local LED closure

The semi-local LED closures (Maltba et al. 2018; Barajas-Solano and Tartakovsky 2016) aim to combine the simplicity and speed of the local closure with (most of) the accuracy of the nonlocal one. One such closure is derived in “Appendix 1.3”,

$$\langle \mathbf{v}'\Pi' \rangle \approx - \int_0^t \mathcal{J}(s; \mathbf{X}, t) \exp((t-s)\mathbf{J}(\mathbf{X}, t))^{\top} \nabla_{\mathbf{X}} \cdot \left(\langle \mathbf{v}'(\mathbf{X}, t) \mathbf{v}'^{\top}(\boldsymbol{\Phi}(s; \mathbf{X}, t), s) \rangle \mathcal{J}^{-1}(s; \mathbf{X}, t) f_{\mathbf{x}}(\mathbf{X}; t) \right) ds, \quad (8)$$

where \mathbf{J} is the Jacobian of the mean-field velocity, $\mathbf{J}(\mathbf{X}, t) \triangleq \partial \langle \mathbf{v}(\mathbf{X}, t) \rangle / \partial \mathbf{X}$. Numerical evaluation of $\boldsymbol{\Phi}(s; \mathbf{X}, t)$ in \mathcal{J} is often computationally expensive. To alleviate this cost, we approximate it by its terminal condition, $\boldsymbol{\Phi}(s; \mathbf{X}, t) \approx \tilde{\boldsymbol{\chi}}$, resulting in the approximate Jacobian determinant $\tilde{\mathcal{J}}$.

Remark 2.1 As an alternative, one can approximate $\boldsymbol{\Phi}(s; \mathbf{X}, t)$ via a higher-order Taylor expansion around the terminal condition (\mathbf{X}, t) . Our numerical experiments in “Numerical experiments” Section reveal that this approximation gives no significant improvement in accuracy, while simultaneously presenting more numerical challenges than the terminal approximation.

Evaluation of the inverse approximate Jacobian determinant $\tilde{\mathcal{J}}^{-1}$ introduces additional challenges that would otherwise be absent if \mathcal{J}^{-1} were known exactly, including, but not limited to, singularities and/or negative values of $\langle \mathbf{v}'\Pi' \rangle(\mathbf{X}, t)$. Many of these issues may be resolved by approximating $\tilde{\mathcal{J}}^{-1}$ with a nonnegative spline. Another option, which entirely circumvents such issues, is to use a simpler but slightly less accurate semi-local closure,

$$\langle \mathbf{v}'\Pi' \rangle \approx - \int_0^t \tilde{\mathcal{J}}(s; \mathbf{X}, t) \exp((t-s)\mathbf{J}(\mathbf{X}, t))^{\top} \nabla_{\mathbf{X}} \cdot \left(\langle \mathbf{v}'(\mathbf{X}, t) \mathbf{v}'^{\top}(\boldsymbol{\Phi}(s; \mathbf{X}, t), s) \rangle f_{\mathbf{x}}(\mathbf{X}; t) \right) ds. \quad (9)$$

It is obtained by approximating $f_{\mathbf{x}}(\boldsymbol{\Phi}(s; \mathbf{X}, t); s) \approx \mathbf{f}_{\mathbf{x}}(\mathbf{X}; t)$ in (7), and amounts to assuming the PDF $f_{\mathbf{x}}$ to be approximately constant over the interval $(t - \tau, t)$.

Data-driven closures

Depending on the model, the integrals in (56), (57), (8), and (9) may or not be analytically computable. If they cannot be solved analytically, their numerical evaluation can be computationally expensive, and prohibitively so in high dimensions. As an alternative to the LED closures presented above, we explore empirical (data-informed) closures of (4). The latter are constructed by representing the unknown term $\langle \mathbf{v}'\Pi' \rangle(\mathbf{X}, t)$ as a differential operator

$$\langle \mathbf{v}'\Pi' \rangle \approx \beta_0(\mathbf{X}, t)f_{\mathbf{x}}(\mathbf{X}; t) + \mathbf{D}(\mathbf{X}, t)\nabla_{\mathbf{x}}f_{\mathbf{x}}(\mathbf{X}; t) + \text{higher-order derivatives}, \tag{10}$$

where \mathbf{D} is a second-order semi-positive definite tensor with elements D_{ij} ($i, j = 1 \dots, N$). The M real-valued coefficients $\boldsymbol{\beta} = \{\beta_0, D_{11}, \dots, D_{NN}, \dots\}$ are learned from post-processed Monte Carlo realizations of (1) by solving the following optimization problem.

For simplicity, we consider a square spatial domain $\mathcal{D} \subset \mathbb{R}^N$ and discretize it with a uniform square mesh $X_n^k = X_n^0 + k\Delta X$ for $k \in \{1, K\}$ in each spatial dimension n . For the temporal domain $(0, T_f]$, we also take the uniform grid $t^l = l\Delta t$ for $l \in \{1, L\}$. Let $\hat{f}_{\mathbf{x}} \in \mathbb{R}^{K \times \dots \times K \times L}$ with entries $f_{\mathbf{x}}^{k_1 \dots k_N l} \equiv \hat{f}_{\mathbf{x}}(X_1^{k_1}, \dots, X_N^{k_N}; t^l)$ denote a joint PDF computed via post-processing of Monte Carlo realizations of (1) with a kernel density estimator. We emphasize that $\hat{f}_{\mathbf{x}}$ is not the converged Monte Carlo solution, f_{MC} , of (1). Instead, $\hat{f}_{\mathbf{x}}$ is computed using only a small fraction of the number of trials needed to compute the converged solution f_{MC} with a prescribed accuracy. Furthermore, we denote by $\hat{\mathcal{L}}$ a suitable numerical discretization (e.g., finite difference, finite volume, total variation regularized differentiation, etc.) of the differential operator $\mathcal{L} \equiv \nabla_{\mathbf{x}} \cdot (\langle \mathbf{v} \rangle f_{\mathbf{x}} + \langle \mathbf{v}'\Pi' \rangle)$ in (4), with $\langle \mathbf{v}'\Pi' \rangle$ given by (10). Finally, we introduce the discretized residual

$$\mathcal{R}_{k_1 \dots k_N l}(\boldsymbol{\beta}) \triangleq \frac{\partial f_{\mathbf{x}}^{k_1 \dots k_N l}}{\partial t} + \hat{\mathcal{L}}(f_{\mathbf{x}}^{k_1 \dots k_N l}; \boldsymbol{\beta});$$

an example of its construction is discussed in detail in “Numerics Section”. The optimal variable coefficient vector $\boldsymbol{\beta}$ is one that satisfies

$$\hat{\boldsymbol{\beta}}(\mathbf{X}, t) = \operatorname{argmin}_{\boldsymbol{\beta}(\mathbf{X}, t)} \left\{ \frac{1}{K^N L} \sum_{k_1=1}^K \dots \sum_{k_N=1}^K \sum_{l=1}^L \|\mathcal{R}_{k_1 \dots k_N l}(\boldsymbol{\beta})\|_2^2 + \lambda \|\boldsymbol{\beta}\|_1^2 \right\}, \tag{11}$$

where $\|\cdot\|_2$ is the L_2 norm, $\|\cdot\|_1$ is the L_1 norm penalty ensuring sparsification of features, and λ is a hyper-parameter iteratively chosen through a geometric sequence

(Brunton et al. 2016; Schaeffer 2017; Bakarji and Tarkovsky 2021).

In order to generalize the optimization problem in (11), one must balance a bias-variance tradeoff when designing a hypothesis class, i.e., choosing the order of the derivatives in (10). A large hypothesis class, represented above by M variable coefficients $\boldsymbol{\beta}$, minimizes bias at the cost of high variance. Such a class is more likely to fit an appropriate operator to the data in $\hat{f}_{\mathbf{x}}$; however, at the cost of possibly making the equations too difficult to manipulate analytically. Choosing too simple of a hypothesis class may prevent generalization by automatically removing hypotheses with high variance, i.e., filtering out noise and outliers in the data. By choosing to include both variable coefficients (as opposed to constants) and the L_1 penalty in (11), we aim to balance the bias-variance tradeoff by respectively increasing the power of the model, and by introducing sparsification, which leads to more interpretable equations.

The hypothesis class is further constrained by the fact that (4) is an equation for the PDF of (1), which means that it must take the form of a master equation. As a consequence of Pawula’s theorem (Risken and Frank 1996; van Kampen 1976), the (finite) Kramers-Moyal (KM) expansion of the PDF master equation does not contain any partial derivatives of order higher than two; otherwise, it could yield a PDF that either is negative or may not integrate to one. Hence, the order of derivatives in (10) is restricted to one.

Numerics

For any of the above closures, we solve (4) through operator splitting by successively considering the two equations

$$\frac{\partial f_{\mathbf{x}}}{\partial t} + \nabla_{\mathbf{x}} \cdot \langle \mathbf{v} \rangle f_{\mathbf{x}} = 0, \tag{12}$$

$$\frac{\partial f_{\mathbf{x}}}{\partial t} + \nabla_{\mathbf{x}} \cdot \langle \mathbf{v}'\Pi' \rangle = 0. \tag{13}$$

The conservative advection equation (12) is solved through the `Clawpack` package (Clawpack Development Team 2019) using a Lax-Wendroff discretization and MC limiter. In the $N = 1$ dimensional setting, (13) is solved via a Crank-Nicolson time discretization. For dimensions higher than one, the modified Craig-Sneyd scheme, presented in ‘t Hout and Mishra (2013), is used with $\theta = 1/2$. This stable alternating direction implicit scheme treats mixed derivatives explicitly while maintaining second-order temporal accuracy. In all cases, (13) is spatially discretized with second-order central differencing. `Clawpack` provides two options for operator splitting, namely Godunov or Strang splitting. However, Strang splitting is not

currently available for adaptive mesh refinement. For the experiment in “[Stochastic NS-LIF model](#)” Section, the classic version (uniform spatial mesh) of `Clawpack` is used with Strang splitting since adaptive mesh refinement is not necessary. For the experiment in “[Stochastic FHN model](#)” Section, we use the adaptive version of `Clawpack` (AMRClaw) with Godunov splitting, which significantly reduces computational costs. Although Godunov splitting is only formally first-order accurate, a convergence study revealed that near second-order temporal accuracy is maintained, i.e., its temporal truncation errors range between $\mathcal{O}(\Delta t^{1.9})$ and $\mathcal{O}(\Delta t^2)$.

The computational domain is taken as $[-L, L]^N$ where $L > 0$ is the smallest number large enough to ensure that the domain truncation with homogeneous boundary conditions generates only negligible errors, by which we mean that if the size of each dimension of the domain is increased 10 fold, the error improves only by 0.01 percent.² For a fixed time t , the error is defined in terms of the Kullback-Leibler (KL) divergence (Mackay 2003),

$$D_{\text{KL}}(f_{\text{MC}}||f_{\hat{x}})(t) \triangleq \int_{\mathbb{R}^N} f_{\text{MC}}(\mathbf{X}; t) \ln \left(\frac{f_{\text{MC}}(\mathbf{X}; t)}{f_{\hat{x}}(\mathbf{X}; t)} \right) d\mathbf{X}.$$

It quantifies the amount of information lost when $f_{\hat{x}}$ is used to approximate the high-resolution Monte Carlo solution f_{MC} , which is treated as a yardstick.

Each Monte Carlo realization is achieved by solving (1) with an adaptive fourth-order Runge-Kutta (RK4) scheme, treating the discretized stochastic fluctuation \mathbf{v}' as a variable coefficient. It is well known that standard Runge-Kutta schemes can be implemented for RODEs; however, they typically do not maintain their traditional orders of accuracy (see for instance Neckel and Rupp 2013; Asai and Kloeden 2016 and the references therein). A convergence study revealed that the standard RK4 scheme maintained third-order accuracy for the experiments in “[Numerical experiments](#)” Section. For each time step in the domain, a large number of Monte Carlo realizations are then post-processed with a Gaussian kernel density estimator to obtain the yardstick PDF f_{MC} on the discretized spatial domain. The exact number of realizations used to compute the yardstick solution is problem dependent. It is determined by a convergence study and is explicitly given for each experiment in “[Numerical experiments](#)” Section. The bandwidth of the kernel density estimator is chosen by Silverman’s rule. In addition to f_{MC} , we also compute the data array $f_{\hat{x}}$ first defined in “[Data-driven closures](#)”

Section. The number of realizations used to compute $f_{\hat{x}}$ typically ranges from 5% to 20% of those used for f_{MC} .

Remark 3.1 Other natural choices for measuring error of PDF solutions are the Wasserstein and total variation distances. Relative L^p norms are typically poor choices for measuring error of PDFs. They often underrepresent the accuracy of closures, sometimes increasing the error by one order of magnitude even if many statistical measures such as mean, variance, skewness, kurtosis, and modality match.

Remark 3.2 Asymptotic “plug-in” estimators of bandwidth, such as Scott’s or Silverman’s rule, are computationally inexpensive. Yet, if the underlying distribution is heavy-tailed or multimodal, computing the yardstick PDF f_{MC} and the data array $f_{\hat{x}}$ may require an unusually large number of realizations. More robust approaches to bandwidth selection include various forms of cross-validation (CV) (e.g., k-fold CV Hastie et al. 2009, one-sided CV Savchuk and Hart 2017, [47], etc.) and total variation regularization [48]. The latter allows for variable bandwidths, but has to solve a bandwidth optimization problem at each time step.

LED closures

The local and semi-local formulations of the LED closures have the same computational complexity, while the non-local formulation is considerably more expensive (Maltba et al. 2018). Evaluation of the cross-covariance term (7) in the latter involves a numerical approximation of the integral, which requires the history of the solution up to the current time step. In the results presented below, the proposed local and semi-local approaches are four orders of magnitude faster than the standard Monte Carlo approach.

Data-driven closures

Numerical solutions of (12) and (13) with semi-local or nonlocal LED closures (7)–(9) may become highly non-trivial, e.g., when $N > 3$ and \mathbf{v}' is spatially dependent. Our numerical experiments reported in “[Numerical Experiments](#)” Section reveal the total computational costs of the data-driven closures to be higher than that of the local and semi-local LED closures. Yet, the data-driven closures are still considerably cheaper than the standard Monte Carlo approach and, in contrast to the LED closures, are relatively easy to implement numerically. This makes the data-driven closures an attractive alternative. Moreover, if the optimization problem (11) is properly constrained and not naively implemented, the bulk computational costs of the data-driven approach lies in generating the (relatively few)

² In “[Stochastic NS-LIF model](#)” Section, the solution to (1) is positive, so the domain is reduced to $(0, L]$ with a homogeneous boundary condition at $X = L$ and a reflective (i.e., zero-flux) boundary condition at $X = 0$.

Monte Carlo realizations needed to assemble \hat{f}_x . The remainder of this section is devoted to solving (11).

Remark 3.3 A naive strategy for solving the optimization problem (11), which involves the discretized variable coefficients $\beta(X_1^{k_1}, \dots, X_N^{k_N}, t^l)$ at each spatiotemporal grid point $(X_1^{k_1}, \dots, X_N^{k_N}, t^l)$, is prohibitively expensive without proper computational resources, particularly memory allocation. The experiment in “Stochastic FHN model” Section, which has a relatively low-dimensional ($N = 2$) phase space, illustrates the high-dimensional nature of this optimization problem. An accurate solution of (12) would require a spatial grid size of at most $\Delta X_1 = \Delta X_2 = 0.01$ at the highest level of mesh refinement. Although this is handled with an adaptive mesh when solving (12) and (13), a uniform mesh is used for the optimization problem (11). For a square domain, a uniform mesh of size $\Delta X_1 = \Delta X_2 = 0.01$ consisting of 1000 elements in each direction would result in 10^6 spatial grid cells. Given that the time step Δt must satisfy the hyperbolic CFL condition, the total number of spatiotemporal grid points is on the order of 10^9 for each of the M coefficients β . The resulting dimension of the optimization problem is $\mathcal{O}(MK^NL) = M10^9$, well beyond the limitations of a personal computer.

To reduce the dimensionality of the optimization problem (11), we use a polynomial basis expansion to represent the decision variables $\beta(\mathbf{X}, t)$. In this setting, it is natural to directly reference the elements of $\beta(\mathbf{X}, t)$ in terms of β_0 and \mathbf{D} from (10). For a polynomial (e.g., Chebyshev, Legendre, etc.) basis $P_q(\cdot)$,

$$\beta_0(\mathbf{X}, t) \triangleq \sum_{q_1=1}^{Q_1} \dots \sum_{q_N=1}^{Q_N} \sum_{r=1}^R \alpha_0^{q_1 \dots q_N r} P_{q_1}(X_1) \dots P_{q_N}(X_N) P_r(t), \tag{14a}$$

$$D_{ij}(\mathbf{X}, t) \triangleq \sum_{q_1=1}^{Q_1} \dots \sum_{q_N=1}^{Q_N} \sum_{r=1}^R \alpha_{ij}^{q_1 \dots q_N r} P_{q_1}(X_1) \dots P_{q_N}(X_N) P_r(t). \tag{14b}$$

The optimization problem (11) is now solved over the expansion coefficients $\alpha_0^{q_1 \dots q_N r}$ and $\alpha_{ij}^{q_1 \dots q_N r} \in \mathbb{R}$ for $i, j \in \{1, N\}$. This step reduces the optimization dimension from $\mathcal{O}(MK^NL)$ to $\mathcal{O}(MQ_1 \dots Q_N R)$, where each $Q_1, \dots, Q_N, R < 10$. For the experiments in “Numerical experiments” Section, the optimization dimension is reduced from $M10^9$ to approximately $M10^1$.

To formulate the regression problem, we take the divergence of (10). Ignoring higher-order derivatives yields

$$\nabla \cdot \langle \mathbf{v}' \Pi' \rangle \approx \sum_{i=1}^N \partial_{X_i} (\beta_0 f_x) + \sum_{i,j=1}^N \partial_{X_i} [D_{ij} (\partial_{X_j} f_x)]. \tag{15}$$

We let $\mathcal{P}_{q_1 \dots q_N r}^{k_1 \dots k_N l} \triangleq P_{q_1}(X_1^{k_1}) \dots P_{q_N}(X_N^{k_N}) P_r(t^l)$ denote the discretized polynomial bases. We also build a target vector $\mathbf{V} \in \mathbb{R}^{K^NL}$ in (12) that contains the advection terms for each spatiotemporal grid point approximated with a Lax-Wendroff discretization and MC limiter. Its elements, denoted by the multi-index $k_1 \dots k_N l$, are given by

$$V^{k_1 \dots k_N l} \triangleq (\partial_t \hat{f}_x + \nabla_x \cdot \langle \mathbf{v} \rangle \hat{f}_x)(X_1^{k_1}, \dots, X_N^{k_N}, t^l).$$

After substituting the discretized version of (14) into (15) and changing the order of summations, for a fixed spatiotemporal grid node with index $k_1 \dots k_N l$, the discretized residual takes the form

$$\begin{aligned} \mathcal{R}_{k_1 \dots k_N l} &\triangleq V^{k_1 \dots k_N l} \\ &+ \sum_{q_1 \dots q_N r} \sum_i \partial_{X_i} \left[\mathcal{P}_{q_1 \dots q_N r}^{k_1 \dots k_N l} f_x^{k_1 \dots k_N l} \right] \alpha_0^{q_1 \dots q_N r} \\ &+ \sum_{i,j} \sum_{q_1 \dots q_N r} \partial_{X_i} \left[\mathcal{P}_{q_1 \dots q_N r}^{k_1 \dots k_N l} (\partial_{X_j} f_x^{k_1 \dots k_N l}) \right] \alpha_{ij}^{q_1 \dots q_N r}. \end{aligned} \tag{16}$$

Rewriting in matrix/vector form gives

$$\mathcal{R} = \mathbf{V} + \mathbf{F}_0 \alpha_0 + \sum_{i,j} \mathbf{F}_{ij} \alpha_{ij}, \tag{17}$$

where the elements of each matrix are given by

$$\mathbf{F}_0(k_1 \dots k_N l, q_1 \dots q_N r) \triangleq \sum_i \partial_{X_i} \left[\mathcal{P}_{q_1 \dots q_N r}^{k_1 \dots k_N l} f_x^{k_1 \dots k_N l} \right], \tag{18}$$

$$\mathbf{F}_{ij}(k_1 \dots k_N l, q_1 \dots q_N r) \triangleq \partial_{X_i} \left[\mathcal{P}_{q_1 \dots q_N r}^{k_1 \dots k_N l} (\partial_{X_j} f_x^{k_1 \dots k_N l}) \right], \tag{19}$$

for all $i, j \in \{1, N\}$. Here, the rows of \mathbf{F}_0 and each \mathbf{F}_{ij} have the multi-index $k_1 \dots k_N l$ and the columns $q_1 \dots q_N r$. We compute all numerical derivatives in (18) and (19) via central finite differencing.

Since the diffusion tensor \mathbf{D} is symmetric, for a fixed index $q_1 \dots q_N r$, $\alpha_{ij}^{q_1 \dots q_N r} = \alpha_{ji}^{q_1 \dots q_N r}$ for all $i, j \in \{1, N\}$. Hence, the residual in (17) is written as

$$\mathcal{R} = \mathbf{V} + \mathbf{F}_0 \alpha_0 + \sum_i \mathbf{F}_{ii} \alpha_{ii} + \sum_{i < j} (\mathbf{F}_{ij} + \mathbf{F}_{ji}) \alpha_{ij}. \tag{20}$$

Then the optimization problem over all polynomial coefficients takes the form

$$\begin{aligned}\hat{\mathbf{A}} &\triangleq [\hat{\alpha}_0, \hat{\alpha}_{11}, \dots, \hat{\alpha}_{NN}, \hat{\alpha}_{12}, \dots, \hat{\alpha}_{1N}, \hat{\alpha}_{23}, \dots, \hat{\alpha}_{(N-1)N}]^\top \\ &= \operatorname{argmin}_{\mathbf{A}} \|\mathcal{R}(\mathbf{A})\|_2^2 + \lambda \|\mathbf{A}\|_1,\end{aligned}\quad (21a)$$

where

$$\begin{aligned}\mathcal{R}(\mathbf{A}) &\triangleq \mathbf{V} + [\mathbf{F}_0, \mathbf{F}_{11}, \dots, \mathbf{F}_{NN}, \mathbf{F}_{12} + \mathbf{F}_{21}, \dots, \mathbf{F}_{1N} \\ &\quad + \mathbf{F}_{N1}, \mathbf{F}_{23} + \mathbf{F}_{32}, \dots, \mathbf{F}_{(N-1)N} + \mathbf{F}_{N(N-1)}] \mathbf{A}.\end{aligned}\quad (21b)$$

We solve (21) by utilizing the lasso function of MATLAB with a 10-fold CV set on each iteration. We use a mean squared error loss function and a relative tolerance of 10^{-4} , i.e., the algorithm terminates when the L^2 norm of successive estimates of the coefficient vector differ by less than 10^{-4} . The extrapolation power of our model is tested by fitting the data for the first half of the time horizon, $[0, T_f/2]$, and analyzing the model performance on the remaining half, $(T_f/2, T_f]$.

Remark 3.4 It may be of interest to consider different variations of CV algorithms that are not considered in this study, including least-angle regression and Lasso implemented with Bayes information criterion. Such algorithms may give different results depending on the choice of the hypothesis class.

Numerical Experiments

We consider two applications of the method of distributions and its associated closures to neuron models, namely, the NS-LIF and FHN models, both driven by SW channel noise. The first model has one state variable ($N = 1$), and the second has two ($N = 2$). SW noise is commonly used in neuroscience, e.g., to model signal transmission (Kang et al. 2013), resonance (Liang et al. 2010), and detection (Liang et al. 2010) in nervous systems; to induce spike death in excitable systems (Guo et al. 2012); and for weak periodic signal detection in the FHN neuron (Yao and Ma 2018).

SW noise $\zeta(t)$ is a bounded and correlated zero-mean non-Gaussian process that is defined by

$$\zeta(t) \triangleq \sigma \sin\left(\sqrt{\frac{2}{\tau}} W_t\right), \quad (22)$$

where $\sigma > 0$ is the finite noise amplitude/strength, $\tau > 0$ is the correlation length, and W_t is a standard Wiener process. It has mean $\langle \zeta(t) \rangle = 0$ for any fixed time t , and covariance function (Ning and Liu 2016; Yang and Ning 2017; van Kampen 1976; Risken and Frank 1996)

$$\begin{aligned}C_\zeta(t, s) &\triangleq \langle \zeta(t) \zeta(s) \rangle \\ &= \frac{\sigma^2}{2} \exp\left(-\frac{t-s}{\tau}\right) \left[1 - \exp\left(-\frac{4s}{\tau}\right)\right], \quad s \leq t.\end{aligned}\quad (23)$$

For the experiments in “Stochastic NS-LIF model” Section, the RODE (1) and all of its corresponding PDF equations are solved up to time $T_f = 50$, while in “Stochastic FSH model” Section, they are solved up to $T_f = 100$.

Stochastic NS-LIF model

LIF neuron models are often used in place of their higher-dimensional counterparts, i.e., models with large N , especially when considering large coupled networks. They describe the basic neuron behavior of slowly charging in the presence of current, spiking after reaching a voltage threshold, and then slowly discharging (Gerstner and Kistler 2002). In order to present an example where all of the proposed closures are in agreement, the LIF model under consideration is not a spiking model, i.e., it lacks a voltage threshold and resetting mechanism for the membrane potential. We do study spiking dynamics in “Stochastic FSH model” Section for which there is a discrepancy between the different closures.

In particular, we consider a (one-dimensional) LIF model that is embedded in the (two-dimensional) stochastic FHN model in the excitable regime.³ The derivation of the embedding is presented in “Numerical experiments” Section of Yamakou et al. (2019). We use their Eq. 4.3, albeit driven by SW noise rather than Gaussian white noise and without a threshold and resetting mechanism, to describe the membrane voltage $x(t) > 0$,

$$\frac{d x}{d t} = \frac{\sigma^2}{2x} - \mu x + \zeta(t), \quad (24)$$

with a given initial condition $x(0) = x_0 \sim |\mathcal{N}(0.1, 0.0125^2)|$ (i.e., a folded Gaussian random variable). In the simulations below, we set $\mu = 0.0312496$ and consider noise amplitudes $\sigma = 0.0444, 0.0666, 0.0888$, and 0.1333 , as in Yamakou et al. (2019). These amplitudes have been appropriately scaled from $\sigma = 0.005, 0.0075, 0.01$, and 0.015 in the FHN model. The correlation lengths vary between $\tau = 0.05, 0.10, 0.15$, and 0.20 .

The RODE (24) is a one-dimensional version of (1) with $v(x(t), t; \zeta(t)) \equiv \sigma^2/(2x) - \mu x + \zeta$. In phase space, $v(X, t; \zeta(t)) \equiv \sigma^2/(2X) - \mu X + \zeta$ is linear in the additive

³ A FHN neuron is considered to be in the excitable regime when starting in the basin of attraction of a unique stable fixed point, a large excursion, or spike, occurs in phase space and then permanently returns to the fixed point (Izhikevich 2007).

SW noise $\xi(t)$. Hence, $\langle v \rangle \equiv \sigma^2/(2X) - \mu X$ and $v' \equiv \xi$. The terminal problem (5) takes the form

$$\frac{d\chi}{ds} = \frac{\sigma^2}{2\chi} - \mu\chi, \quad \chi(t) = X, \tag{25}$$

and admits an analytical solution, yielding a closed-form expression for the associated flow $\chi(s) = \Phi(s; X, t)$.

Local LED closure

For the problem under consideration, equations (56) and (57), which define the drift velocity \mathcal{V} and the diffusion coefficient \mathcal{D} in the local LED closure (6), reduce to

$$\mathcal{V}(X, t) \equiv 0 \quad \text{and} \quad \mathcal{D}(X, t) = \int_0^t \mathcal{J}(s; X, t) C_\xi(t, s) \, ds, \tag{26}$$

where

$$\begin{aligned} \mathcal{J}(s; X, t) &= \exp\left(-\int_s^t \frac{\partial}{\partial \chi} \langle v(\chi(r), r) \rangle \, dr\right) \\ &= \exp\left(\int_s^t \left[\frac{\sigma^2}{2\Phi^2(r; X, t)} + \mu\right] \, dr\right). \end{aligned} \tag{27}$$

The integral in \mathcal{J} can be evaluated analytically, but this is generally not the case. To be consistent with the approximations that are usually needed, we approximate $\Phi(s; X, t)$ in \mathcal{J} by its terminal condition X , resulting in the approximate Jacobian

$$\begin{aligned} \tilde{\mathcal{J}}(s; X, t) &= \exp\left(\int_s^t \left[\frac{\sigma^2}{2X^2} + \mu\right] \, dr\right) \\ &= \exp\left[(t-s)\left(\frac{\sigma^2}{2X^2} + \mu\right)\right]. \end{aligned} \tag{28}$$

With this local LED closure, the single-point PDF $f_x(X; t)$ of the random membrane voltage $x(t)$ satisfies the PDF equation

$$\frac{\partial f_x}{\partial t} + \frac{\partial}{\partial X} \left[\left(\frac{\sigma^2}{2X} - \mu X\right) f_x \right] = \frac{\partial}{\partial X} \left(\tilde{\mathcal{D}}(X, t) \frac{\partial f_x}{\partial X} \right), \tag{29}$$

where the diffusion coefficient $\tilde{\mathcal{D}}$ is computed upon replacing \mathcal{J} with $\tilde{\mathcal{J}}$ in (26) and analytically evaluating the resulting integral.

Remark 4.1 As mentioned in Remark 2.1, a higher-order Taylor expansion around the terminal condition (X, t) may be used as an alternative approximation for $\Phi(s; X, t)$ in (27). Our numerical experiments with this approximation showed no improvement in accuracy, but simultaneously induced additional numerical challenges: the resulting linear approximation of the Jacobian determinant

\mathcal{J} behaves like $\exp[(t-s)\mathcal{O}(X^{-2})]$, as does $\tilde{\mathcal{J}}$ in (28), but has singularities at points other than $X = 0$.

Semi-local LED closure

For the problem under consideration, the semi-local closure (9) reduces to

$$\langle v' \Pi' \rangle(X, t) \approx -D_2(t) \frac{\partial f_x}{\partial X}, \tag{30}$$

where the diffusion coefficient $D_2(t)$ is computed analytically,

$$\begin{aligned} D_2(t) &= \int_0^t \tilde{\mathcal{J}}(s; X, t) \exp\left((t-s) \frac{d \langle v(X, t) \rangle}{d X}\right) C_\xi(t, s) \, ds \\ &= \int_0^t C_\xi(t, s) \, ds. \end{aligned} \tag{31}$$

The resulting PDF equation for $f_x(X; t)$ is identical to (29), except that now $D_2(t)$ is used in place of $\tilde{\mathcal{D}}(X, t)$.⁴ A solution to this equation, $f_x(X; t)$, is shown in Fig. 1a for several combinations of the meta-parameters σ and τ characterizing the strength and correlation length of the colored noise $\xi(t)$, respectively. For the stochastic NS-LIF model (24), the PDFs of membrane voltage, $f_x(X; t)$, computed via the PDF equation with the nonlocal and local closures are visually indistinguishable from the PDF with the semi-local closure shown in Fig. 1a. Figure 1b gives a direct comparison of the PDF computed with the semi-local closure and the yardstick PDF, $f_{\text{MC}}(X; t)$, estimated via high-resolution Monte Carlo simulations. The latter required $N_{\text{MC}} \approx 1.5 \times 10^5$ realizations to converge. The results in Fig. 1b are displayed for the worst case scenario, which occurs for the largest choices of σ and τ and at early times, as seen by the error plot in Fig. 2.

Figure 2 provides a quantitative assessment of the accuracy of the semi-local closure. It displays the temporal evolution of the KL divergence $D_{\text{KL}}(f_{\text{MC}} || f_x)$ between $f_x(X; t)$ and $f_{\text{MC}}(X; t)$ for the noise strength $\sigma = 0.044$ and 0.133. (The errors associated with the intermediate values of $\sigma = 0.066$ and 0.088 follow the same pattern and, hence, are omitted from Fig. 2). The KL divergence increases with σ and τ , as to be expected from the theoretic considerations, since the LED-based closures can be thought of as leading-order terms in perturbation expansions in the powers of σ and τ . Regardless of the choice of the meta-parameter values, $D_{\text{KL}}(f_{\text{MC}} || f_x)$ generally decreases with time t . That is because the membrane voltage PDF

⁴ We found the semi-local closure (8) to provide only minor improvements in accuracy relative to the closure (9), while being significantly more involved computationally. Consequently, only the numerical results for the semi-local closure (9) are presented therein.

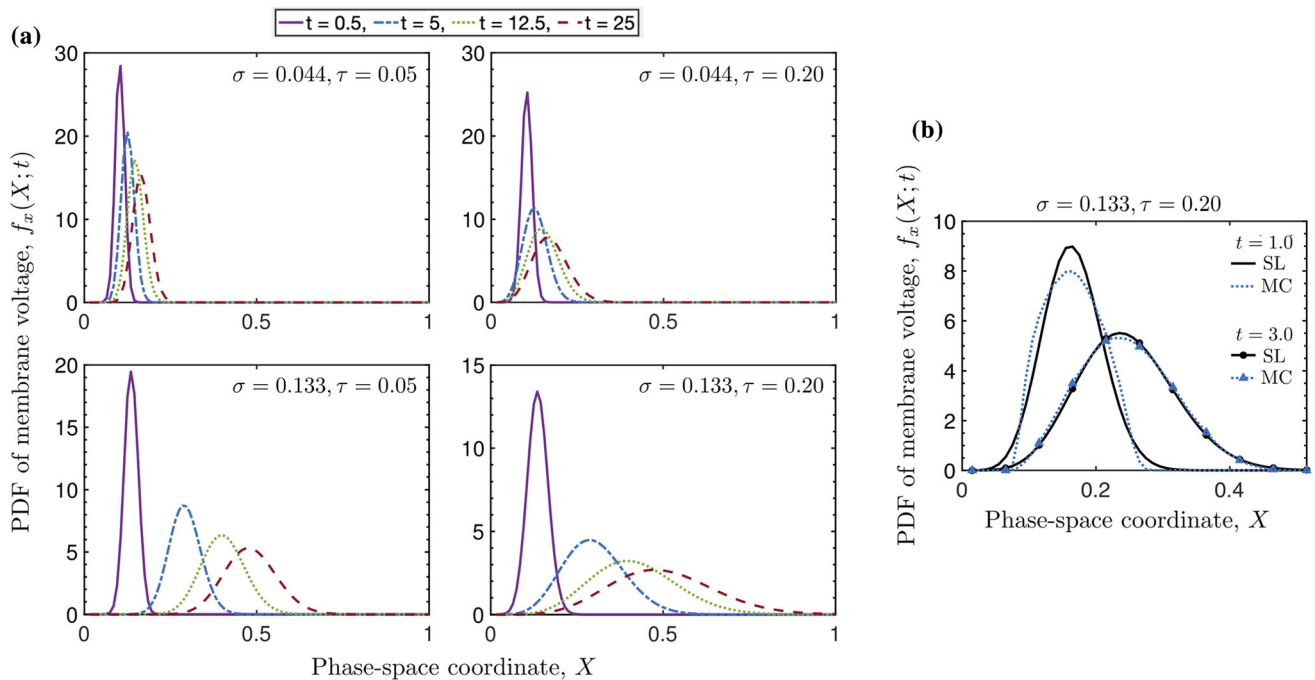



Fig. 1 (a) PDF $f_x(X; t)$ of the membrane voltage $x(t)$ in (24), computed via the PDF equation with the semi-local closure, for several combinations of the strength (σ) and correlation length (τ) of

colored noise $\zeta(t)$. (b) Direct comparison of the PDF $f_x(X; t)$ computed with the semi-local closure and the yardstick PDF $f_{\text{MC}}(X; t)$ for $\sigma = 0.133$ and $\tau = 0.20$ at times $t = 1$ and 3

Fig. 2 Temporal evolution of the KL divergence, $D_{\text{KL}}(f_{\text{MC}}||f_x)$, between the membrane voltage PDF computed with the semi-local closure, $f_x(X; t)$, and its Monte Carlo counterpart, $f_{\text{MC}}(X; t)$, for several combinations of the strength (σ) and correlation length (τ) of colored noise $\zeta(t)$.



approaches its nearly Gaussian steady state, wherein the PDF and its derivatives with respect to X are almost constant over the correlation-length intervals $(t - \tau, t)$, an approximation that underpins both the local and semi-local closures. For all choices of σ and τ and for all time t , the nonlocal, semi-local, and local closures yield PDFs $f_x(X; t)$ within a KL divergence of 1.6×10^{-4} of one another, with the largest error occurring for the largest values of σ and τ .⁵ The local and semi-local approaches are four orders of magnitude faster than the Monte Carlo approach.

⁵ Although the accuracy of the local and semi-local closures is nearly identical for the stochastic NS-LIF model (24), this is generally not the case, as seen in “Stochastic FSH model” Section.

Data-driven closure

Since the diffusion coefficient for the semi-local closure, D_2 , depends only on time, see (31), we consider spatially independent variable coefficients $\beta \equiv \beta(t)$ in the data-driven closure (14), significantly reducing the dimensionality of the optimization problem (11)⁶. Furthermore, since $D_2(t)$ is a monotonically increasing function that quickly reaches an asymptote (Fig. 3), we also investigate a data-driven closure with constant coefficients β , eliminating the need for a polynomial expansion and further simplifying the optimization problem. In both cases, the one-dimensional version of (10) gives rise to the hypothesis

⁶ To be concrete, we represented $\beta_i(t)$ ($i = 1, 2$) with the first ten Legendre polynomials.


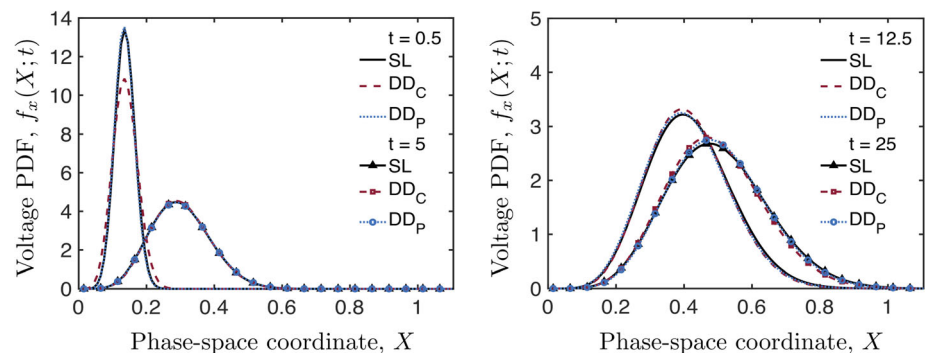


Fig. 3 Diffusion coefficients associated with the semi-local closure, $D_2(t)$, the data-driven closure with constant coefficients, $\hat{\beta}_2$, and the data-driven closure with a temporal Legendre basis expansion, $\hat{\beta}_2(t)$. We present the cases for which the KL divergence $D_{\text{KL}}(f_{\text{MC}}||f_x)$ of

the data-driven closure with constant $\hat{\beta}$ is highest. In all cases, the data-driven coefficients approximated with a Legendre basis expansion agree well with the other closures

Fig. 4 Temporal snapshots (at times $t = 0.5, 5, 12.5$, and 25 ; for $\sigma = 0.133$ and $\tau = 0.20$) of the membrane voltage PDFs $f_x(X; t)$ alternatively computed with the semi-local closure (SL), the data-driven closure with constant coefficients (DD_C), and the data-driven closure with a temporal Legendre basis expansion (DD_P)



$$\frac{\partial \langle v' \Pi' \rangle}{\partial X} \approx \beta_1 \frac{\partial f_x}{\partial X} + \beta_2 \frac{\partial^2 f_x}{\partial X^2}, \quad (32)$$

and either constant or time-dependent coefficients $\beta = (\beta_1, \beta_2)^\top$ are found by sparse regression (11) in which the hyper-parameter λ is chosen to minimize the mean squared error.

Regardless of whether the coefficients β are treated as constant or time-dependent, the sparse regression yields $\beta_1 \equiv 0$, so that β_2 becomes an approximation of the diffusion coefficient $D_2(t)$; the optimal estimates of both constant and time-varying β_2 , denoted by $\hat{\cdot}$, are compared with $D_2(t)$ in Fig. 3. For small values of σ and τ (Fig. 3a), $D_2(t)$ quickly approaches its asymptote, rendering its approximation with the constant $\hat{\beta}_2$ accurate. Meanwhile, the time-dependent approximation $\hat{\beta}_2(t)$ suffers from oscillations around the asymptote, with a noticeable overshoot near $t = 1$. This discrepancy becomes less pronounced as the number of polynomial bases is increased. As σ and τ increase (Fig. 3b), the rate at which $D_2(t)$ reaches its asymptote decreases. This causes the constant approximation $\hat{\beta}_2$ to slightly underestimate $D_2(t)$ because it has to “average out” low levels of diffusion at early times.

However, this slower transition benefits the time-dependent approximation $\hat{\beta}_2(t)$, allowing for fewer polynomial bases and less pronounced oscillations and overshoot.

Figure 4 reveals that, while the semi-local closure outperforms the data-driven closure (it is visually indistinguishable from f_{MC}), the data-driven closures perform remarkably well. For large values of σ and τ , allowing $\hat{\beta}_2$ to vary in time improves the accuracy of the data-driven closure at early times, but gives only minor improvements at later times. To a certain extent, the under-estimation of $D_2(t)$ by the constant $\hat{\beta}_2$ at later times is balanced by over-estimation of $D_2(t)$ at early times. However, the discrepancy does increase at later extrapolated times, as the constant $\hat{\beta}_2$ leads to the data-driven PDF $f_x(X; t)$ being under-diffused. This comparison is provided for the largest values of σ and τ considered in our study ($\sigma = 0.133$ and $\tau = 0.20$).

Figure 5 elaborates upon this visual observation by presenting the temporal evolution of the KL divergence $D_{\text{KL}}(f_{\text{MC}}||f_x)$ between the “exact” solution $f_{\text{MC}}(X; t)$ and its approximations $f_x(X; t)$ obtained with the data-driven closure. We illustrate the performance of the least


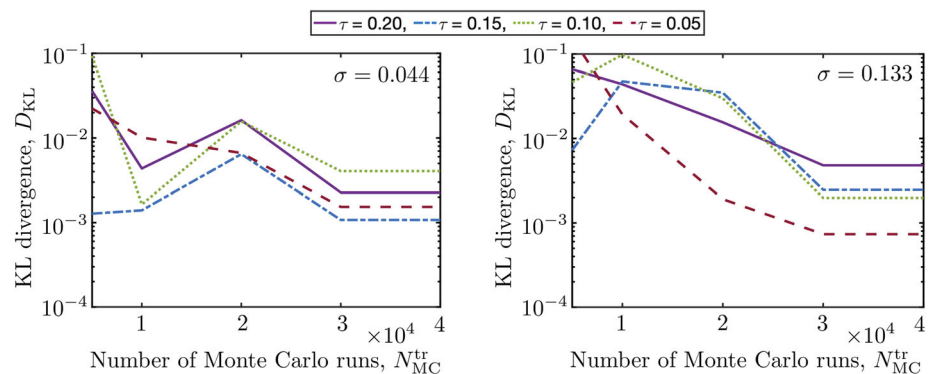


Fig. 5 Temporal evolution of the KL divergence, $D_{\text{KL}}(f_{\text{MC}}||f_x)$, between the membrane voltage PDF computed with the data-driven closure, $f_x(X;t)$, and its Monte Carlo counterpart, $f_{\text{MC}}(X;t)$, for selected combinations of the strength (σ) and correlation length (τ) of

colored noise $\zeta(t)$. The data array \hat{f}_x used in the data-driven closure is computed with $N_{\text{MC}}^{\text{tr}} = 3 \times 10^4$ Monte Carlo realizations

Fig. 6 Dependence of the KL divergence $D_{\text{KL}}(f_{\text{MC}}||f_x)$, averaged over all temporal grid nodes, on the number of Monte Carlo runs used in the sparse regression, $N_{\text{MC}}^{\text{tr}}$.



computationally expensive data-driven closure, i.e., the one that relies on the constant coefficients $\hat{\beta} \in \mathbb{R}^2$. As σ and τ decrease, the discrepancy between the semi-local closure and the data-driven closure increases, with the largest difference in their KL divergences occurring for $\sigma = 0.044$ and $\tau = 0.05$ at later times. At first glance, this may seem contradictory to Figure 3, but one must consider that the diffusion coefficients decrease with σ and τ , and that σ also influences advection as it is included in the drift term in (24). When σ and τ are small, the resulting PDF $f_x(X;t)$ advects and diffuses at a slow rate (see Fig. 1), resulting in a spikier PDF. Hence, in terms of relative entropy, i.e., the KL divergence, $f_x(X;t)$ is more sensitive to small changes in the diffusion coefficient. Unlike in the LED theory, the data-driven closure does not guarantee that $f_x(X;t)$ should improve as σ and τ decrease. Regardless of these observations, the data-driven closure remains accurate when compared to the semi-local closure.

The residual \mathcal{R} in the optimization problem (11) uses the data array \hat{f}_x computed from $N_{\text{MC}}^{\text{tr}}$ Monte Carlo realizations. Hence, the optimal values of β and, thus, the PDF $f_x(X;t)$ computed with (12), (13) and (32) are affected by the choice of $N_{\text{MC}}^{\text{tr}}$. Figure 6 sheds light on this

dependence by exhibiting the KL divergence $D_{\text{KL}}(f_{\text{MC}}||f_x)$, averaged over all temporal grid nodes, as function of $N_{\text{MC}}^{\text{tr}}$. For all choices of σ and τ considered, the optimal solution $\hat{\beta}$ to (11) needs $N_{\text{MC}}^{\text{tr}} \approx 3 \times 10^4$ Monte Carlo trials to stabilize, and the accuracy of the data-driven closure does not improve for $N_{\text{MC}}^{\text{tr}} > 3 \times 10^4$. This is less than 20% of the number of realizations, N_{MC} , needed to compute f_{MC} . The data-driven results presented in Figs. 3, 4, and 5 are obtained with $N_{\text{MC}}^{\text{tr}} = 3 \times 10^4$ Monte Carlo trials.


Stochastic FHN model

The FHN neuron is controlled by two state variables, the membrane potential $x_1(t)$ and the recovery variable $x_2(t)$ that restores a neuron's resting state. In the excitable regime, its dynamics is described by

$$\begin{bmatrix} \dot{x}_1 \\ \dot{x}_2 \end{bmatrix} = \begin{bmatrix} x_1 - x_1^3/3 - x_2 + I \\ \epsilon(x_1 + \alpha - \beta x_2) + \zeta \end{bmatrix} \quad (33)$$

where SW channel noise $\zeta(t)$ is defined by (22); I is the external input current; ϵ is a small timescale parameter that

Fig. 7 The coefficients of the semi-local closure, D_4 and D_5 , and the data-driven closure, $\hat{\beta}_3$ and $\hat{\beta}_4$. For $\sigma = 0.05$ and $\tau = 0.1$, the optimization algorithm sets $\hat{\beta}_3 = 0$ while $D_4 = \mathcal{O}(10^{-7})$, which is near zero when compared with $\hat{\beta}_4$ and D_5 .



ensures that the dynamics of x_1 is much faster than x_2 ; and α and β are model parameters. In the simulations reported below, we set $I = 0.265$, $\epsilon = 0.08$, $\alpha = 0.7$ and $\beta = 0.75$, in accordance with (Yamakou et al. 2019). With these parameter values, the deterministic version of (33), i.e., in the absence of the noise ξ , is in the excitable regime⁷ and has a unique stable fixed point at $(x_1, x_2) = (-1.00125, -0.401665)$. In the stochastic setting, the channel noise $\xi(t)$ may perturb the trajectory away from the fixed point, causing repeated firing. The frequency of this repeated firing increases with the noise amplitude σ (Yamakou et al. 2019). In our simulations we consider noise amplitudes of $\sigma = 0.01, 0.05$, and 0.2 , and correlation lengths $\tau = 0.01, 0.1$, and 5.0 . The initial state $\mathbf{x}(0) = \mathbf{x}_0$ is an uncorrelated joint Gaussian whose mean is perturbed from the stable fixed point:

$$\mathbf{x}_0 \sim \mathcal{N}\left(\begin{bmatrix} -1.00125 \\ -0.411665 \end{bmatrix}, \begin{bmatrix} 0.01 & 0 \\ 0 & 0.01 \end{bmatrix}\right).$$

The approximate Jacobian determinant $\tilde{\mathcal{J}}$ associated with the FHN model (33) needed for both the local and semi-local LED closures is given by

$$\tilde{\mathcal{J}}(s; \mathbf{X}, t) = \exp(-(t-s)(1 - \epsilon\beta - X_1^2)).$$

Local LED closure

For the problem under consideration, equations (56) and (57) with $\mathcal{J} \approx \tilde{\mathcal{J}}$ yield analytical expressions for the drift velocity \mathcal{V} and diffusion coefficient \mathcal{D} ,

$$\mathcal{V}(\mathbf{X}, t) \equiv \begin{bmatrix} 0 \\ 0 \end{bmatrix}, \tag{34a}$$

$$\begin{aligned} \mathcal{D}(\mathbf{X}, t) &\triangleq \int_0^t \tilde{\mathcal{J}}(s; \mathbf{X}, t) \langle \mathbf{v}'(\mathbf{X}, t) \mathbf{v}'^\top(\Phi(s; \mathbf{X}, t), s) \rangle \, ds \\ &= D_3(X_1, t) \begin{bmatrix} 0 & 0 \\ 0 & 1 \end{bmatrix}, \end{aligned} \tag{34b}$$

where

$$D_3(X_1, t) = \int_0^t \tilde{\mathcal{J}}(s; \mathbf{X}, t) C_\xi(t, s) \, ds. \tag{34c}$$

The integral in (34c) is computed analytically. The PDF equation (58) for the joint density $f_{\mathbf{x}}(\mathbf{X}; t)$ simplifies to

$$\frac{\partial f_{\mathbf{x}}}{\partial t} + \nabla_{\mathbf{X}} \cdot \left(\begin{bmatrix} X_1 - X_1^3/3 - X_2 + I \\ \epsilon(X_1 + \alpha - \beta X_2) \end{bmatrix} f_{\mathbf{x}} \right) = D_3 \frac{\partial^2 f_{\mathbf{x}}}{\partial X_2^2}. \tag{35}$$

Semi-local LED closure

The semi-local closure (9) takes the form

$$\begin{aligned} \langle \mathbf{v}'\Pi' \rangle(\mathbf{X}, t) &= - \int_0^t \tilde{\mathcal{J}}(s; \mathbf{X}, t) C_\xi(t, s) \exp((t-s)\mathbf{J}(\mathbf{X}, t))^\top \begin{bmatrix} 0 & 0 \\ 0 & 1 \end{bmatrix} \nabla_{\mathbf{X}} f_{\mathbf{x}}(\mathbf{X}; t) \, ds. \end{aligned} \tag{36}$$

The integrals in (36) are computed exactly with a symbolic computation software Mathematica, giving rise to the PDF equation

$$\begin{aligned} \frac{\partial f_{\mathbf{x}}}{\partial t} + \nabla_{\mathbf{X}} \cdot \left(\begin{bmatrix} X_1 - X_1^3/3 - X_2 + I \\ \epsilon(X_1 + \alpha - \beta X_2) \end{bmatrix} f_{\mathbf{x}} \right) \\ = \frac{\partial}{\partial X_1} \left(D_4 \frac{\partial f_{\mathbf{x}}}{\partial X_2} \right) + D_5 \frac{\partial^2 f_{\mathbf{x}}}{\partial X_2^2}, \end{aligned} \tag{37}$$

where $D_4(X_1, t)$ and $D_5(X_1, t)$ are known functions whose explicit expressions are omitted for brevity, and whose graphical representation is provided in Fig. 7. (The X_1 -dependence of these coefficients is insignificant, and $D_4(\cdot, t)$ and $D_5(\cdot, t)$ are plotted for $X_1 = 0$.) Fig. 8 shows

⁷ A neuron is considered to be in the excitable regime when starting in the basin of attraction of a unique stable fixed point, a large excursion, or spike, occurs in the phase space and then permanently returns to the fixed point (Izhikevich 2007).




Fig. 8 Temporal snapshots of the joint PDF $f_{\mathbf{x}}(\mathbf{X}; t)$ of the random state variables $x_1(t)$ and $x_2(t)$ in the FHN model (33). The dynamics of $f_{\mathbf{x}}(\mathbf{X}; t)$ is governed by the PDF equation (37), for $\sigma = 0.2$ and $\tau = 0.1$.

temporal snapshots of the joint PDF $f_{\mathbf{x}}(\mathbf{X}; \cdot)$ computed with (37). The shape of this PDF reveals that the ensemble mean $\langle \mathbf{x}(t) \rangle$ in and of itself does not accurately capture the dynamics of (33).

Figure 9 depicts the temporal evolution of the KL divergence $D_{\text{KL}}(f_{\text{MC}} \| f_{\mathbf{x}})$. It serves to ascertain the accuracy of the semi-local closure that underpins the PDF equation (37). As expected, the accuracy increases as the correlation length τ decreases. As the noise amplitude σ decreases, the different choices for τ become less impactful as the errors are more tightly clustered. Although not shown here, we found that, for all choices of σ , the local closure agrees with the semi-local closure within a KL divergence of 10^{-4} for $\tau = 0.01$. For $\tau = 0.1$, the KL divergence of the local closure is between 0.1 and 0.5 more than that of the semi-local closure; and for $\tau = 5.0$, it is up to a full order of magnitude more than that of the semi-local closure.

Data-driven closure

Guided by the results reported in “Data-driven closure” Section, we focus on constant coefficients $\boldsymbol{\beta}$ so that the data-driven closure (10) gives rise to a hypothesis

$$\begin{aligned} \nabla_{\mathbf{X}} \cdot \langle \mathbf{v}' \Pi' \rangle = & \beta_1 \left(\frac{\partial f_{\mathbf{x}}}{\partial X_1} + \frac{\partial f_{\mathbf{x}}}{\partial X_2} \right) + \beta_2 \frac{\partial^2 f_{\mathbf{x}}}{\partial X_1^2} \\ & + \beta_3 \frac{\partial^2 f_{\mathbf{x}}}{\partial X_1 \partial X_2} + \beta_4 \frac{\partial^2 f_{\mathbf{x}}}{\partial X_2^2}. \end{aligned} \quad (38)$$

The sparse regression optimization (11) identifies $\hat{\beta}_1 = \hat{\beta}_2 = 0$, so that the resulting PDF equation (4) matches the PDF equation (37) obtained with the semi-local closure, with the constants $\hat{\beta}_3$ and $\hat{\beta}_4$ serving as empirical representations of $D_4(X_1, t)$ and $D_5(X_1, t)$, respectively. The values of $\hat{\beta}_3$ and $\hat{\beta}_4$ are shown in Fig. 7 for two combinations of σ and τ . These values slightly underestimate the asymptotes of the semi-local diffusion coefficients


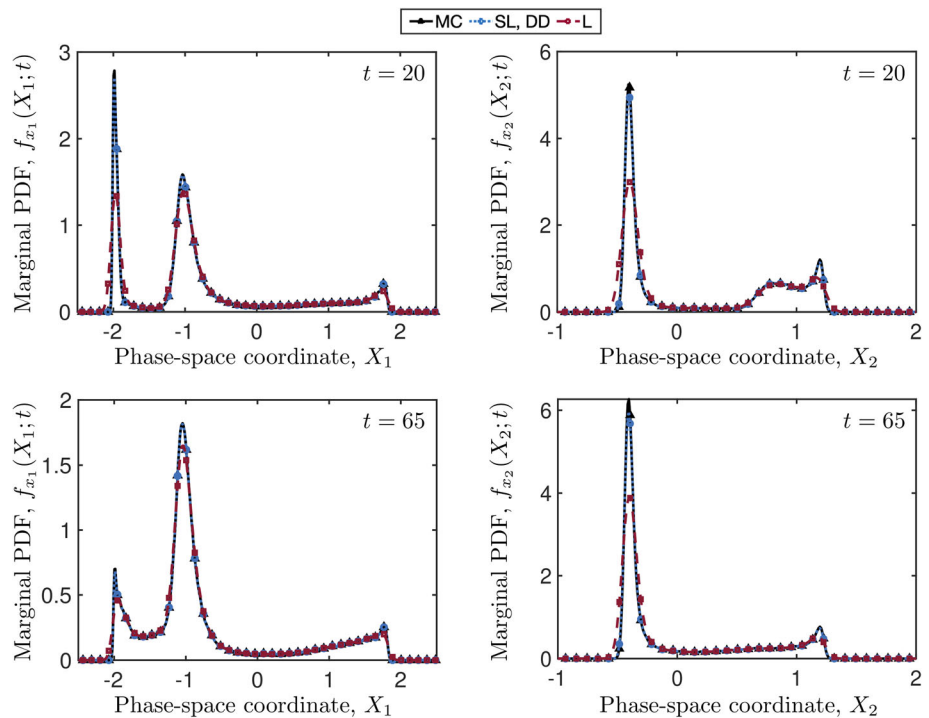


Fig. 9 Temporal evolution of the KL divergence $D_{\text{KL}}(f_{\text{MC}}||f_x)$ between the yardstick Monte Carlo solution $f_{\text{MC}}(\mathbf{X}; t)$ of (33) and the PDF $f_x(\mathbf{X}; t)$ computed, alternatively, via (37) and (4) with the data-driven closure (38)

Fig. 10 Snapshots (at times $t = 20$ and 65) of the marginal PDFs $f_{x_1}(X_1; t)$ (left column) and $f_{x_2}(X_2; t)$ (right column) alternatively computed with the local (L), semi-local (SL) and data-driven (DD) closures and with Monte Carlo simulations (MC), for the stochastic FHN neuron with $\sigma = 0.05$ and $\tau = 0.1$.



$D_4(\cdot, t)$ and $D_5(\cdot, t)$ due to “averaging out” of low levels of diffusion at early times.

This discrepancy gives rise to the relatively small difference in predictions of $f_x(\mathbf{X}; t)$, as quantified by the KL divergence $D_{\text{KL}}(f_{\text{MC}}||f_x)$ in Fig. 9. The semi-local closure outperforms the data-driven closure, but just barely. The close agreement between the PDF predictions based on these two closures is remarkably robust, holding for all choices of σ and τ considered. Although not shown here, both closures are considerably more accurate than the local closure for larger τ , as discussed in “Semi-local LED closure” Section.

The close agreement between predictions of the joint PDF $f_x(\mathbf{X}; t)$ based on the semi-local and data-driven

closures translates into the close agreement between predictions of the corresponding marginal PDFs $f_{x_1}(X_1; t)$ and $f_{x_2}(X_2; t)$ (Fig. 10). The semi-local and data-driven closures accurately capture all features of the Monte Carlo solution. While capturing the coarse dynamics of the Monte Carlo solution, the local closure is too diffusive to accurately capture the dynamics with sharp peaks. Since the PDFs associated with the semi-local and data-driven closures are visually indistinguishable, we include them as the same line-plot.

Finally, we note that an optimal solution $\hat{\beta}(N_{\text{MC}}^{\text{tr}})$ to the sparse regression problem (11) needs $N_{\text{MC}}^{\text{tr}}$ Monte Carlo trials to stabilize for all choices of σ and τ . The accuracy of




Fig. 11 Dependence of the KL divergence $D_{KL}(f_{MC}||f_x)$, averaged over all temporal grid nodes, on the number of Monte Carlo runs used in the sparse regression, N_{MC}^{tr} .

the data-driven closure does not improve for $N_{MC}^{tr} > 4 \times 10^4$; hence, $N_{MC}^{tr} = 4 \times 10^4$ was used in the simulation results reported in Figs. 7, 8, 9, and 10. This number is approximately 20% of the number of trials needed to compute the Monte Carlo solution f_{MC} (i.e., $N_{MC} \approx 2 \times 10^5$).

Mutual Information

Information transmission from a spiking neuron occurs via spike trains, i.e., a sequence of action potentials. Observing a spike train can be thought of as gaining information about the sensory stimulus. Neuroscience regularly employs information theory to formalize this notion on an absolute scale so that quantifiable statements can be made about information transmission rates of neurons (Rieke 1997). Mutual information (MI) between two continuous random variables x_i and x_j is given in terms of the KL divergence of their joint PDF, $f_{x_i x_j}(X_i, X_j)$, and the product of their marginals, $f_{x_i}(X_i)$ and $f_{x_j}(X_j)$:

$$I(x_i; x_j) \triangleq \int \int f_{x_i x_j}(X_i, X_j; t) \ln \left(\frac{f_{x_i x_j}(X_i, X_j; t)}{f_{x_i}(X_i; t) f_{x_j}(X_j; t)} \right) d X_i d X_j = D_{KL}(f_{x_i x_j} || f_{x_i} \otimes f_{x_j});$$

it is symmetric and nonnegative (Mackay 2003). While either of the semi-local and data-driven PDF solutions from “Semi-local LED closure” and “Data-driven closure” Sections can be used to calculate MI between the membrane potential x_1 and the recovery variable x_2 in the FHN neuron, the dynamical system (33) must be modified in order to calculate MI between x_1 and the noise source ξ . Letting $\zeta(t) = \sigma \cos(\sqrt{2/\tau} W_t)$, it follows via Itô’s formula (Øksendal 2003, Theorem 4.1.2) that ξ and ζ satisfy the Itô SDEs

$$\begin{aligned} d \xi_t &= -\frac{1}{\tau} \xi_t d t + \sqrt{\frac{2}{\tau}} \zeta_t d W_t, & \xi(0) &= 0, \\ d \zeta_t &= -\frac{1}{\tau} \zeta_t d t - \sqrt{\frac{2}{\tau}} \xi_t d W_t, & \zeta(0) &= 0. \end{aligned} \tag{39}$$

Adding these two SDEs to (33) yields an exact four-dimensional Fokker-Planck equation. For any fixed time t , its solution is a joint PDF $f_{x_1 x_2 \xi \zeta}(X_1, X_2, \Xi, Z; t)$ for the random processes $x_1(t)$, $x_2(t)$, $\xi(t)$, and $\zeta(t)$. Then, the instantaneous (point-wise in time) MI between the noise source and the membrane potential is found by calculating (via numerical quadrature) the joint PDF $f_{x_1 \xi}(X_1, \Xi; t)$ and the marginal PDFs $f_{x_1}(X_1; t)$ and $f_{\xi}(\Xi; t)$. This process involves solving a high-dimensional PDE and then numerically integrating out phase variables, which can be computationally expensive. In the spirit of the data-driven PDF method, a better approach is to directly learn reduced-order equations for the PDF(s) of interest (e.g., $f_{x_1 \xi}(X_1, \Xi; t)$) from data. Such an approach is presented in Brennan and Venturi (2018) for ODEs with random initial conditions, but can easily be generalized for SDEs driven by standard Wiener processes.

Figure 12 displays the temporal evolution of MI between the various FHN states (including noise) for $\sigma = 0.05$ and $\tau = 0.1$. The MI initially increases in all cases: when the system begins to evolve, a portion of the joint PDF enters into a pseudo period due to the probability of a spike, and more information is gained about the coupling of states. However, there is a stark contrast in the evolution of $I(x_1; x_2)$ when compared to $I(x_1; \xi)$ and $I(x_2; \xi)$. The MI $I(x_1; x_2)$ oscillates according to the pseudo-periodic behavior of the joint PDF (Fig. 8), i.e., the neuron’s spikes. The MI between the membrane potential $x_1(t)$ and the recovery variable (inhibitory response) $x_2(t)$ is gained as the neuron spikes, but decreases as the neuron returns to the vicinity of its resting state, as one would expect (Rieke 1997). Unlike $I(x_1; x_2)$, both $I(x_1; \xi)$ and $I(x_2; \xi)$ are nearly constant in time after a slight initial increase. Hence, apart




Fig. 12 Temporal evolution of MI, I , between the different FHN neuron states and of the total correlation, C , between all three states. The PDF equation (4) is closed with the data-driven closure, and the parameter values are set to $\sigma = 0.05$ and $\tau = 0.01$.

from the first spike, observing the membrane potential/recovery variable for a fixed time does not provide any more or less information about the noise source than any of the other temporal observations. The converse is also true since MI is symmetric. This behavior is attributed to the fact that the SDE for $\xi(t)$ is independent from the states $x_1(t)$ and $x_2(t)$. The noise ξ can change the dynamics of x_1 and x_2 by perturbing the neuron from rest and causing it to fire; however, x_1 and x_2 have no effect on how the noise is generated.

While MI is limited to measuring information transfer between two states/variables, its many generalizations allow one to do the same for N -many variables x_1, \dots, x_N . One such generalization is total correlation (Vejnarová 1999; Watanabe 1960). Like MI, it is defined in terms of the KL divergence of the variables' joint PDF, $f_{x_1 \dots x_N}$, and the product of their marginals. Their total correlation is defined as

$$C(x_1, \dots, x_N) = D_{\text{KL}}(f_{x_1 \dots x_N} \| f_{x_1} \otimes \dots \otimes f_{x_N}).$$

Figure 12 shows the total correlation between x_1 , x_2 , and ξ for the FHN neuron as an example of another information-theoretic quantity that can be calculated via the PDF method.

Remark 4.2 The PDF method in “Method of distributions” Section can be generalized to derive an equation for a joint PDF of system states at multiple times. This enables the calculation of information transfer rates between states at different times, as opposed to the instantaneous information transmission. For example, instead of calculating total correlation between the FHN states x_1 and x_2 and the noise source ξ at any time t , one can compute the total correlation between these states at any two times t_1 and t_2 . This would lead to the doubling of the number of dimensions in the PDF equation from three to six, and require deployment of numerical strategies for solutions of high-

dimensional PDEs such as tensor-train methods (Boelens et al. 2018).

Remark 4.3 Apart from computing information-theoretic metrics, the PDF method can be used to calculate PDFs for first exit/spike times and interspike-intervals. For example, the FHN neuron spikes (almost surely) when its path crosses the line $X_1 = 0$ (Yamakou et al. 2019). Defining the first exit/spike time to be $\tau_1 = \inf\{s > 0 : x_1(s) = 0\}$, its distribution can be calculated in terms of the cumulative distribution function (CDF) $F_{\mathbf{x}}(X_1, X_2; t)$. When studying first spiking times, it is common to consider deterministic initial conditions in (1). This setting corresponds to the initial condition for the PDF equation (4) being the Dirac delta function, which introduces significant numerical challenges [56]. This can be circumvented by deploying the CDF method (see Tartakovsky and Gremaud 2015 and the references therein). It would result in a PDE for the CDF of system states, for which the deterministic initial condition in (1) translates into the Heaviside function as the initial CDF.

Conclusions

Nonlinear dynamical systems provide a rich framework for understanding complex neuronal processes. Their stochastic analogues allow one to study the effects of externally imposed and/or internally generated noise on neuronal dynamics. They can also be used to quantify parametric uncertainty in underlying models. Noise sources in ODE-based neuron models are often treated as Gaussian white noise due to its convenient mathematical properties, one being that the joint PDF of system states satisfies exactly the Fokker-Planck equation. However, Gaussian white noise is not a biologically plausible noise source for many neuronal systems, for which colored noise characterized by a short correlation length is a more realistic representation.

Monte Carlo methods are often employed to calculate PDF solutions of stochastic and random ODEs, as well as information-theoretic quantities and distributions of various spiking statistics. They are easily implementable, yet computationally demanding. Efficient methods for general Langevin-type systems with colored noise have been studied extensively; however, most are not appropriate for neuron models. Methods utilizing only the first few moments of system states typically underperform because low-order moments can spend considerable time in a low probability state, as illustrated in Fig. 8 by the ensemble mean. Additionally, many methods that do characterize a full PDF solution, such as polynomial chaos expansions and stochastic finite elements, suffer from the curse of dimensionality when applied to systems with short

correlation lengths. The method of distributions, comprised of PDF and CDF methods, is the exception; this study utilizes the former to discover Fokker-Planck-type equations for stochastic neuron models. Such equations describe PDF dynamics of neuronal states, and in general, the states of general Langevin-type systems. While the PDF equations resulting from the method of distributions are exact for ODEs driven by white noise (i.e., SDEs), they are generally unclosed and require approximations in the presence of colored noise. We proposed two approaches for constructing such closures: nonlocal large eddy diffusivity (LED) closures and a data-driven closure relying on statistical inference; we apply each to study the stochastic NS-LIF and FHN neurons.

Fully nonlocal LED closures result in derived integro-partial differential PDF equations that are difficult to discretize numerically and computationally demanding to solve unless the underlying RODEs exhibit weak nonlinearities and are low-dimensional, e.g., the NS-LIF model. For more complicated models, e.g., the FHN neuron, localization approximations are needed to render the PDF equation numerically computable. Full localization yields “easy-to-solve” linear partial-differential PDF equations, but introduces significant approximation error unless the noise intensity and correlation length are small. Partial (semi-) localization provides an attractive trade-off between accuracy and computational complexity; however, this is likely to be problem dependent. Regardless, LED closures reduce the computational cost by several orders of magnitude when compared to the standard Monte Carlo approach, with the local and semi-local LED closures being more efficient than the nonlocal one.

Unlike LED closures, the proposed data-driven closure needs only a few Monte Carlo trials and kernel density estimation to learn terms in the PDF equation from a dictionary of possible differentiable operators. The relevant derivative features and their coefficients are found by formulating an optimization problem via sparse regression, i.e., Lasso. If the number of equations in a neuron model is greater than one, the dimensionality of the regression problem is inherently high when considering arbitrary variable coefficients. It can be significantly reduced by approximating the coefficients with a polynomial basis expansion, which is sufficiently accurate for the models considered in this study. While more computationally expensive than the local and semi-local LED closures, the data-driven closure is easier to implement and still less expensive than the standard Monte Carlo approach. Its accuracy is comparable to the semi-local LED closure.

The method of distributions is a computationally efficient approach for finding PDF solutions to stochastic neuron models driven by colored noise. Post-processing of these solutions yields distributions of spiking statistics or

information-theoretic quantities such as mutual information. For simple models, e.g., the NS-LIF neuron, all of the proposed closures yield consistent results and remain accurate over a wide range of noise intensities and correlation lengths. For models with more complicated dynamics, e.g., the FHN model and other two- or three-dimensional models such as Morris-Lecar, the semi-local LED and data-driven closures are likely to be in agreement and maintain good accuracy as long as $\sigma\tau$, the product of the noise intensity and correlation length, is not very large relative to the system’s scale. For high-dimensional models, such as Hodgkin-Huxley (HH), the numerical discretization of PDF equations resulting from LED closures may be difficult, and the data-driven closure is likely to give better results.

In a follow-up study, we will investigate both the HH neuron as well systems of coupled neurons, all driven by colored noise. In the latter case, we aim to use the method of distributions to determine how model complexity (i.e., using LIF versus FHN versus HH neurons, etc.) affects information content and transmission in networks of neurons. Such a study will require deployment of numerical strategies for solutions of high-dimensional PDEs such as tensor-train methods (Boelens et al. 2018; Dektor et al. 2021; Rodgers et al. 2021).

Appendix 1. Derivation of PDF Equations

Appendix 1.1. Regularization Argument

Let us define $\Pi_\epsilon(\mathbf{X}, t)$, a regularized version of $\Pi(\mathbf{X}, t)$ in (2), as

$$\begin{aligned}\Pi_\epsilon(\mathbf{X}, t) &\triangleq (\eta_\epsilon \star \Pi)(\mathbf{X}, t) \triangleq \int_{\mathbb{R}^N} \eta_\epsilon(\mathbf{X} - \mathbf{Y}) \Pi(\mathbf{Y}, t) \, d\mathbf{Y} \\ &= \eta_\epsilon(\mathbf{X} - \mathbf{x}),\end{aligned}\quad (40)$$

where the last equality holds by the definition of $\Pi(\mathbf{Y}, t)$ and the sifting property of the Dirac distribution. The standard positive mollifier $\eta_\epsilon \in \mathcal{C}_0^\infty(\mathbb{R}^N)$ satisfies the conditions of symmetry, $\eta_\epsilon(\mathbf{X} - \mathbf{x}) = \eta_\epsilon(\mathbf{x} - \mathbf{X})$, and scaling

$$\begin{aligned}\eta_\epsilon(\mathbf{Y}) &\triangleq \frac{\epsilon^{-N}}{\int \eta \, d\mathbf{Y}} \eta\left(\frac{\mathbf{Y}}{\epsilon}\right), \quad \text{where } \eta(\mathbf{Y}) \\ &\triangleq \begin{cases} \exp\left(\frac{1}{|\mathbf{Y}|^2 - 1}\right) & \text{if } |\mathbf{Y}| < 1 \\ 0 & \text{if } |\mathbf{Y}| \geq 1. \end{cases}\end{aligned}\quad (41)$$

Following the arguments from Evans (2010), one can show that Π_ϵ is a smooth approximation of Π . Let $\phi(\mathbf{X}, t) \in \mathcal{C}_c^1(\mathbb{R}^N \times [0, \infty))$. It follows from (40) that

$$\begin{aligned} \mathcal{I} &\triangleq \int_0^\infty \int_{\mathbb{R}^N} \Pi_\epsilon(\mathbf{X}, t) \frac{\partial \phi}{\partial t}(\mathbf{X}, t) \, d\mathbf{X} \, dt \\ &= \int_0^\infty \int_{\mathbb{R}^N} \eta_\epsilon(\mathbf{X} - \mathbf{x}) \frac{\partial \phi}{\partial t}(\mathbf{X}, t) \, d\mathbf{X} \, dt. \end{aligned} \tag{42}$$

Integrating by parts in t and applying the sifting property gives

$$\begin{aligned} \mathcal{I} &= \int_0^\infty \int_{\mathbb{R}^N} \dot{\eta}_\epsilon(\mathbf{X} - \mathbf{x}) \mathbf{v}(\mathbf{x}, t, \omega) \phi(\mathbf{X}, t) \, d\mathbf{X} \, dt \\ &\quad - \int_{\mathbb{R}^N} \eta_\epsilon(\mathbf{X} - \mathbf{x}_0) \phi(\mathbf{X}, 0) \, d\mathbf{X}, \\ &= \int_0^\infty \int_{\mathbb{R}^N} \int_{\mathbb{R}^N} \dot{\eta}_\epsilon(\mathbf{X} - \mathbf{Y}) \mathbf{v}(\mathbf{Y}, t, \omega) \Pi(\mathbf{Y}, t) \phi(\mathbf{X}, t) \\ &\quad d\mathbf{Y} \, d\mathbf{X} \, dt \\ &\quad - \int_{\mathbb{R}^N} \Pi_\epsilon(\mathbf{X}, 0) \phi(\mathbf{X}, 0) \, d\mathbf{X}, \end{aligned}$$

where $\dot{\eta}_\epsilon(\cdot)$ is the derivative of $\eta_\epsilon(\cdot)$. According to the Gauss-Ostrogradsky theorem in \mathbf{X} ,

$$\begin{aligned} \mathcal{I} &= - \int_0^\infty \int_{\mathbb{R}^N} (\eta_\epsilon \star \mathbf{v} \Pi)(\mathbf{X}, t) \nabla_{\mathbf{X}} \phi(\mathbf{X}, t) \, d\mathbf{X} \, dt \\ &\quad - \int_{\mathbb{R}^N} \Pi_\epsilon(\mathbf{X}, 0) \phi(\mathbf{X}, 0) \, d\mathbf{X}. \end{aligned} \tag{43}$$

It follows from (43) and (42) that, for any $\phi \in \mathcal{C}_c^1(\mathbb{R}^N \times [0, \infty))$,

$$\begin{aligned} \int_0^\infty \int_{\mathbb{R}^N} \Pi_\epsilon \frac{\partial \phi}{\partial t} \, d\mathbf{X} \, dt + \int_0^\infty \int_{\mathbb{R}^N} (\eta_\epsilon \star \mathbf{v} \Pi) \nabla_{\mathbf{X}} \phi \, d\mathbf{X} \, dt \\ + \int_{\mathbb{R}^N} \Pi_\epsilon(\mathbf{X}, 0) \phi(\mathbf{X}, 0) \, d\mathbf{X} = 0. \end{aligned}$$

By standard arguments, taking the limit $\epsilon \rightarrow 0$ gives

$$\begin{aligned} \int_0^\infty \int_{\mathbb{R}^N} \Pi \frac{\partial \phi}{\partial t} \, d\mathbf{X} \, dt + \int_0^\infty \int_{\mathbb{R}^N} (\mathbf{v} \Pi) \nabla_{\mathbf{X}} \phi \, d\mathbf{X} \, dt \\ + \int_{\mathbb{R}^N} \Pi(\mathbf{X}, 0) \phi(\mathbf{X}, 0) \, d\mathbf{X} = 0, \end{aligned} \tag{44}$$

Hence, Π is the distributional solution to (3).

Appendix 1.2. LED closure derivation

Let L denote the linear operator

$$L\Pi \triangleq \frac{\partial \Pi}{\partial t} + \nabla_{\mathbf{X}} \cdot (\langle \mathbf{v} \rangle \Pi)$$

together with the initial and boundary conditions

$$\Pi(\mathbf{X}, 0) = \delta(\mathbf{x}^0 - \mathbf{X}) \quad \text{and} \quad \lim_{|\mathbf{X}| \rightarrow \infty} \Pi(\mathbf{X}, t) = 0. \tag{45}$$

Let \hat{L} denote the adjoint of L , i.e.,

$$\hat{L}\Psi \triangleq - \frac{\partial \Psi}{\partial s} - \langle \mathbf{v} \rangle \cdot \nabla_{\mathbf{Y}} \Psi$$

with vanishing boundary condition at infinity and vanishing terminal condition $\Psi(\mathbf{Y}, t) = 0$. Finally, let $G \equiv G(\mathbf{X}, t; \mathbf{Y}, s)$ be the Green's function of \hat{L} ; it is defined as a solution of

$$\hat{L}G = \delta(\mathbf{X} - \mathbf{Y}) \delta(t - s), \tag{46}$$

with vanishing boundary condition at infinity and terminal condition $G(\mathbf{X}, t; \mathbf{Y}, t) = 0$. We use the method of characteristics to compute this Green's function. The characteristics of (46), $\boldsymbol{\varphi} \equiv \boldsymbol{\varphi}(r)$ with $s \leq r \leq t$, satisfy

$$\frac{d\boldsymbol{\varphi}}{dr} = \langle \mathbf{v}(\boldsymbol{\varphi}, r) \rangle, \quad r \in (s, t), \tag{47}$$

subject to the initial condition

$$\boldsymbol{\varphi}(s) = \mathbf{Y}. \tag{48}$$

It induces the associated flow $\boldsymbol{\Phi}$ as $\boldsymbol{\varphi}(r) \equiv \boldsymbol{\Phi}(r; \mathbf{Y}, s)$. Along these characteristics, (46) reduces to

$$\frac{d}{dr} G(\mathbf{X}, t; \boldsymbol{\varphi}, r) = -\delta(t - r) \delta(\mathbf{X} - \boldsymbol{\varphi}), \quad r \in (s, t);$$

with $G(\mathbf{X}, t; \boldsymbol{\Phi}(t; \mathbf{Y}, s), t) = 0$. Its solution gives the Green's function,

$$G(\mathbf{X}, t; \mathbf{Y}, s) = \mathcal{H}(t - s) \delta(\mathbf{X} - \boldsymbol{\Phi}(t; \mathbf{Y}, s)), \tag{49}$$

where \mathcal{H} is the Heaviside function.

Elementary manipulations lead to the equation for Π' ,

$$L\Pi' = -\nabla_{\mathbf{X}} \cdot (\mathbf{v}' \Pi - \langle \mathbf{v}' \Pi' \rangle), \tag{50}$$

with homogeneous initial conditions and vanishing boundary conditions. Rewriting (50) in terms of \mathbf{Y} and s , multiplying by $G(\mathbf{X}, t; \mathbf{Y}, s)$, and integrating by parts, we obtain

$$\begin{aligned} \int_0^t \int_{\mathbb{R}^N} GL\Pi' \, d\mathbf{Y} \, ds = - \int_{\mathbb{R}^N} G(\mathbf{X}, t; \mathbf{Y}, 0) \Pi'(\mathbf{Y}, 0) \, d\mathbf{Y} \\ + \int_0^t \int_{\mathbb{R}^N} \Pi' \hat{L}G \, d\mathbf{Y} \, ds, \end{aligned}$$

where the additional boundary terms cancel because of the terminal condition on G and the vanishing boundary condition (45) at infinity. Accounting for (46) and (50),

$$\begin{aligned} \Pi'(\mathbf{X}, t) &= \int_{\mathbb{R}^N} G(\mathbf{X}, t; \mathbf{Y}, 0) \Pi'(\mathbf{Y}, 0) \, d\mathbf{Y} \\ &\quad + \int_0^t \int_{\mathbb{R}^N} G(\mathbf{X}, t; \mathbf{Y}, s) \nabla_{\mathbf{Y}} \\ &\quad \cdot \left(\langle \mathbf{v}'(\mathbf{Y}, s) \Pi'(\mathbf{Y}, s) \rangle - \mathbf{v}'(\mathbf{Y}, s) \Pi(\mathbf{Y}, s) \right) \, d\mathbf{Y} \, ds. \end{aligned}$$

The first integral vanishes if the initial condition on \mathbf{x}^0 is deterministic. An exact, albeit unclosed, equation for the stochastic flux $\langle \mathbf{v}' \Pi' \rangle$ is obtained by multiplying the previous relation with $\mathbf{v}'(\mathbf{X}, t)$ and taking the ensemble mean,

$$\begin{aligned} \langle \mathbf{v}' \Pi' \rangle(\mathbf{X}, t) &= \int_{\mathbb{R}^N} G(\mathbf{X}, t; \mathbf{Y}, 0) \langle \mathbf{v}'(\mathbf{X}, t) \Pi'(\mathbf{Y}, 0) \rangle \, d\mathbf{Y} \\ &\quad - \int_0^t \int_{\mathbb{R}^N} G(\mathbf{X}, t; \mathbf{Y}, s) \nabla_{\mathbf{Y}} \\ &\quad \cdot \langle \mathbf{v}'(\mathbf{X}, t) \mathbf{v}'^T(\mathbf{Y}, s) \Pi(\mathbf{Y}, s) \rangle \, d\mathbf{Y} \, ds. \end{aligned}$$

The LED closure is constructed by setting the first integral to zero, resulting in a second-order (in $\sigma\tau$) computable relationship

$$\begin{aligned} \langle \mathbf{v}' \Pi' \rangle(\mathbf{X}, t) &\approx - \int_0^t \int_{\mathbb{R}^N} G(\mathbf{X}, t; \mathbf{Y}, s) \nabla_{\mathbf{Y}} \\ &\quad \cdot \left(\langle \mathbf{v}'(\mathbf{X}, t) \mathbf{v}'^T(\mathbf{Y}, s) \rangle f_{\mathbf{y}}(\mathbf{Y}, s) \right) \, d\mathbf{Y} \, ds. \end{aligned} \tag{51}$$

Substituting G from (49) and setting

$$q(\mathbf{X}, \mathbf{Y}, t, s) \triangleq \nabla_{\mathbf{Y}} \cdot \left(\langle \mathbf{v}'(\mathbf{X}, t) \mathbf{v}'^T(\mathbf{Y}, s) \rangle f_{\mathbf{y}}(\mathbf{Y}, s) \right),$$

we obtain

$$\begin{aligned} \langle \mathbf{v}' \Pi' \rangle(\mathbf{X}, t) &\approx - \int_0^t \int_{\mathbb{R}^N} \delta(\mathbf{X} - \Phi(t; \mathbf{Y}, s)) q(\mathbf{X}, \mathbf{Y}, t, s) \, d\mathbf{Y} \, ds \\ &\approx - \int_0^t \int_{\mathbb{R}^N} \delta(\mathbf{X} - \varphi) q(\mathbf{X}, \Phi(s; \varphi, t), t, s) \left| \frac{\partial \Phi(s; \varphi, t)}{\partial \varphi} \right| \, d\varphi \, ds. \end{aligned}$$

The Jacobian determinant in this expression is made explicit by invoking the Liouville-Ostrogradsky lemma. It is reproduced and proved below because we were not able to find an appropriate reference.

Lemma 1 (Liouville-Ostrogradsky). *Consider an initial value problem*

$$\frac{d\boldsymbol{\chi}}{ds} = \mathbf{g}(\boldsymbol{\chi}, s), \quad \boldsymbol{\chi}(t) = \boldsymbol{\xi}, \tag{52}$$

where $\mathbf{g} : \mathbb{R}^N \times \mathbb{R} \rightarrow \mathbb{R}^N$ is a given deterministic function, and $\boldsymbol{\xi}$ is a possibly random initial condition. Both \mathbf{g} and $\boldsymbol{\xi}$ are such that, for any realization of $\boldsymbol{\xi}$, the above system admits a unique solution for the range of time under

consideration. Then, if Φ is the corresponding flow, i.e., $\boldsymbol{\chi}(s) \equiv \Phi(s; \boldsymbol{\xi}, t)$, we have

$$\left| \frac{\partial \Phi(s; \boldsymbol{\xi}, t)}{\partial \boldsymbol{\xi}} \right| = \exp \left(- \int_s^t \nabla_{\boldsymbol{\chi}} \cdot \mathbf{g}(\boldsymbol{\chi}(r), r) \, dr \right).$$

Proof Following the arguments of “Method of distributions” Section, we show that f , the PDF of $\boldsymbol{\chi}$, satisfies

$$\frac{\partial f}{\partial t} + \nabla_{\boldsymbol{\chi}} \cdot (\mathbf{g}f) = 0.$$

By construction, the restriction of f on characteristic curves, $\mathcal{P}(s) = f(\boldsymbol{\chi}(s); s)$, satisfies

$$\frac{d\mathcal{P}}{ds} = -\mathcal{P} \nabla_{\boldsymbol{\chi}} \cdot \mathbf{g}.$$

Hence,

$$\mathcal{P}(s) = \mathcal{P}(t) \exp \left(- \int_t^s \nabla_{\boldsymbol{\chi}} \cdot \mathbf{g}(\boldsymbol{\chi}(r), r) \, dr \right).$$

On the other hand, treating the flow Φ as a one-to-one change of variables,

$$\mathcal{P}(s) = \left| \frac{\partial \Phi^{-1}}{\partial \boldsymbol{\chi}} \right| \mathcal{P}(t).$$

Combining these two expressions, and taking the inverse completes the proof. \square

We apply the Liouville-Ostrogradsky lemma to $\boldsymbol{\chi}$ and Φ defined, respectively, as the solution of and the flow corresponding to (52) with $\mathbf{g}(\boldsymbol{\chi}, s) = \langle \mathbf{v}(\boldsymbol{\chi}, s) \rangle$ and $\boldsymbol{\xi} = \mathbf{X}$. This yields

$$\begin{aligned} \langle \mathbf{v}' \Pi' \rangle(\mathbf{X}, t) &\approx - \int_0^t \mathcal{J}(s; \mathbf{X}, t) \nabla_{\Phi} \\ &\quad \cdot \left(\langle \mathbf{v}'(\mathbf{X}, t) \mathbf{v}'^T(\Phi(s; \mathbf{X}, t), s) \rangle f_{\mathbf{x}}(\Phi(s; \mathbf{X}, t); s) \right) \, ds, \end{aligned} \tag{53}$$

where

$$\mathcal{J}(s; \mathbf{X}, t) = \exp \left(- \int_s^t \nabla_{\boldsymbol{\chi}} \cdot \langle \mathbf{v}(\boldsymbol{\chi}(r), r) \rangle \, dr \right). \tag{54}$$

Substituting (53) into (4) gives an integro-differential equation for $f_{\mathbf{x}}$. Its numerical solution poses significant challenges as the resulting problem is local neither in time nor in space.

The localization inherent in the classical LED theory (Kraichnan 1987) assumes that $f_{\mathbf{x}}$ and its spatial derivatives are approximately constant over the correlation-length interval $(t - \tau, t)$. This gives rise to approximations $f_{\mathbf{x}}(\Phi(s; \mathbf{X}, t); s) \approx f_{\mathbf{x}}(\mathbf{X}; t)$ and $\nabla_{\Phi} f_{\mathbf{x}}(\Phi(s; \mathbf{X}, t); s) \approx \nabla_{\mathbf{x}} f_{\mathbf{x}}(\mathbf{X}; t)$, so that

$$\langle \mathbf{v}'\Pi' \rangle(\mathbf{X}, t) \approx \mathcal{V}(\mathbf{X}, t)f_{\mathbf{x}}(\mathbf{X}; t) - \mathcal{D}(\mathbf{X}, t)\nabla_{\mathbf{x}}f_{\mathbf{x}}(\mathbf{X}; t), \tag{55}$$

where \mathcal{V} and \mathcal{D} are the LED drift velocity vector and diffusion tensor,

$$\mathcal{V}(\mathbf{X}, t) \triangleq - \int_0^t \mathcal{J}(s; \mathbf{X}, t) \langle \mathbf{v}'(\mathbf{X}, t) \nabla_{\Phi} \cdot \mathbf{v}'^{\top}(\Phi(s; \mathbf{X}, t), s) \rangle \, ds, \tag{56}$$

$$\mathcal{D}(\mathbf{X}, t) \triangleq \int_0^t \mathcal{J}(s; \mathbf{X}, t) \langle \mathbf{v}'(\mathbf{X}, t) \mathbf{v}'^{\top}(\Phi(s; \mathbf{X}, t), s) \rangle \, ds. \tag{57}$$

Substituting (55)–(57) into (4) results in the classical LED equation for the joint PDF $f_{\mathbf{x}}$:

$$\frac{\partial f_{\mathbf{x}}}{\partial t} + \nabla_{\mathbf{x}} \cdot (\langle \mathbf{v} \rangle + \mathcal{V})f_{\mathbf{x}} = \nabla_{\mathbf{x}} \cdot (\mathcal{D} \nabla_{\mathbf{x}} f_{\mathbf{x}}). \tag{58}$$

The shortcomings of the above LED closure—most importantly, its limited validity to short correlation lengths—are investigated in Barajas-Solano and Tartakovsky (2016); Maltba et al. (2018); Ye et al. (2004), among many other studies.

Appendix 1.3 Semi-local LED closure

This closure is based on the mean-field advection problem

$$\frac{\partial f_{\mathbf{x}}}{\partial s} + \nabla_{\mathbf{x}} \cdot (\langle \mathbf{v} \rangle f_{\mathbf{x}}) = 0, \quad s < t \tag{59}$$

and is equivalent to

$$f_{\mathbf{x}}(\Phi(s; \mathbf{X}, t); s) = \mathcal{J}^{-1}(s; \mathbf{X}, t)f_{\mathbf{x}}(\mathbf{X}; t). \tag{60}$$

Equation (60) gives a local approximation for $f_{\mathbf{x}}$; substituting it into (53) yields the semi-local closure

$$\langle \mathbf{v}'\Pi' \rangle(\mathbf{X}, t) \approx - \int_0^t \mathcal{J}(s; \mathbf{X}, t) \nabla_{\Phi} \cdot \left(\langle \mathbf{v}'(\mathbf{X}, t) \mathbf{v}'^{\top}(\Phi(s; \mathbf{X}, t), s) \rangle \mathcal{J}^{-1}(s; \mathbf{X}, t)f_{\mathbf{x}}(\mathbf{X}; t) \right) \, ds. \tag{61}$$

The PDF $f_{\mathbf{x}}$ in (61) has been localized from $(\Phi(s; \mathbf{X}, t); s)$ to $(\mathbf{X}; t)$, but the operator $\nabla_{\Phi} \cdot$ has not. We now approximate ∇_{Φ} in terms of $\nabla_{\mathbf{x}}$,

$$\nabla_{\mathbf{x}} = \Psi^{\top} \nabla_{\Phi} \quad \text{with} \quad \Psi_{ij} \triangleq \frac{\partial \Phi_i}{\partial \mathbf{X}_j}(s; \mathbf{X}, t). \tag{62}$$

Elementary calculus shows that the transpose of the sensitivity matrix of the flow, $\Psi(s; \mathbf{X}, t)$, satisfies

$$\frac{d \Psi}{d s} = \mathbf{J}\Psi, \quad \Psi(t; \mathbf{X}, t) = \mathbb{I},$$

where $\mathbf{J}(\mathbf{Y}, s) \triangleq \partial \langle \mathbf{v}(\mathbf{Y}, s) \rangle / \partial \mathbf{Y}$ is the Jacobian of the mean-

field velocity, and \mathbb{I} is the $N \times N$ identity matrix. Consequently, Ψ can be expressed as

$$\begin{aligned} \Psi(s; \mathbf{X}, t) &= \text{OE}[\mathbf{J}](s) \\ &\triangleq \mathbb{I} + \int_t^s \mathbf{J}(\Phi(s_1; \mathbf{X}, t), s_1) \, ds_1 \\ &\quad + \int_t^s \int_t^{s_1} \mathbf{J}(\Phi(s_1; \mathbf{X}, t), s_1) \mathbf{J}(\Phi(s_2; \mathbf{X}, t), s_2) \, ds_1 \, ds_2 + \dots, \end{aligned}$$

where OE is the ordered exponential. Using the approximation $\mathbf{J}(\Phi(s; \mathbf{X}, t), s) \approx \mathbf{J}(\mathbf{X}, t)$, the ordered exponential simplifies to the matrix exponential. Hence, we obtain

$$\Psi \approx \exp((s - t)\mathbf{J}(\mathbf{X}, t)),$$

and, accounting for (62),

$$\nabla_{\Phi} \approx \exp((t - s)\mathbf{J}(\mathbf{X}, t))^{\top} \nabla_{\mathbf{x}}. \tag{63}$$

Substituting (63) into (61), we obtain the approximate semi-local closure (8).

Acknowledgements This work was supported in part by Air Force Office of Scientific Research Under Award Number FA9550-21-1-0381, National Science Foundation Graduate Research Fellowship under Grant No. DGE 1752814, and by UC Berkeley SURF L&S Program. A special thanks to Michael R. DeWeese for his useful comments regarding the manuscript.

References

Alzubaidi H, Shardlow T (2014) Improved simulation techniques for first exit time of neural diffusion models. *Comm Stat Simul Comput* 43(10):2508–2520

Asai Y, Kloeden PE (2016) Numerical schemes for random odes with affine noise. *Numer Algor* 72(12):155–171. <https://doi.org/10.1007/s11075-015-0038-y>

Bakarji J, Tartakovsky DM (2021) Data-driven discovery of coarse-grained equations. *J. Comput. Phys.* 434:110219. <https://doi.org/10.1016/j.jcp.2021.110219>

Barajas-Solano DA, Tartakovsky AM (2016) Probabilistic density function method for nonlinear dynamical systems driven by colored noise. *Phys Rev E* 93:052121-1-052121-13

Boelens AMP, Venturi D, Tartakovsky DM (2018) Parallel tensor methods for high-dimensional linear PDEs. *J Comput Phys* 375(12):519–539. <https://doi.org/10.1016/j.jcp.2018.08.057>

Boelens AMP, Venturi D, Tartakovsky DM (2020) Tensor methods for the Boltzmann-BGK equation. *J Comput Phys.* <https://doi.org/10.1016/j.jcp.2020.109744>

Brennan C, Venturi D (2018) Data-driven closures for stochastic dynamical systems. *J Comput Phys* 372:281–298. <https://doi.org/10.1016/j.jcp.2018.06.038>

Brunton SL, Proctor JL, Kunz JN, Bialek W (2016) Discovering governing equations from data by sparse identification of nonlinear dynamical systems. *Proc Natl Acad Sci USA* 113(15):3932–3937

Cho H, Venturi D, Karniadakis G (2013) Adaptive discontinuous Galerkin method for response-excitation PDF equations. *SIAM J Sci Comput* 35:B890–B911

Clawpack Development Team (2019) Clawpack software, version 5.6.1. <http://www.clawpack.org>

- Dektor A, Rodgers A, Venturi D (2021) Rank-adaptive tensor methods for high-dimensional nonlinear PDEs. *J Sci Comput* 88(36):1–27. <https://doi.org/10.1007/s10915-021-01539-3>
- Evans L (2010) *Partial Differential Equations*, 2nd edn. AMS, Providence
- FitzHugh R (1961) Impulses and physiological states in theoretical models of nerve membrane. *Biophys J* 1(6):445–466
- Gerstner W, Kistler WM (2002) *Spiking neuron models: single neurons, populations, plasticity*. Cambridge University Press, Cambridge
- Gerstner W, Kistler WM, Naud R, Paninski L (2014) *Neuronal dynamics*. Cambridge University Press, Cambridge
- Giles M (2008) Multilevel Monte Carlo path simulation. *Oper Res* 56(3):607–617
- Greenwood PE, Ward LM (2016) *Stochastic neuron models*. Springer, Cham
- Guo D (2011) Inhibition of rhythmic spiking by colored noise in neural systems. *Cogn Neurodyn* 5(3):293–300
- Guo W, Du LC, Mei DC (2012) Coherence and spike death induced by bounded noise and delayed feedback in an excitable system. *Eur Phys J B* 85(6):182
- Haskell E, Nykamp D, Tranchina D (2001) A population density method for large-scale modeling of neuronal networks with realistic synaptic kinetics. *Neurocomputing* 38:627–632
- Hastie T, Tibshirani R, Friedman J (2009) *The elements of statistical learning*, 2nd edn. Springer, New York. <https://doi.org/10.1007/978-0-387-84858-7>
- Hodgkin AL, Huxley AF (1952) A quantitative description of membrane current and its application to conduction and excitation in nerve. *J Physiol* 117(4):500–544
- Hout in't KJ, Mishra C, (2013) Stability of ADI schemes for multidimensional diffusion equations with mixed derivative terms. *Appl Numer Math* 74:83–94. <https://doi.org/10.1016/j.apnum.2013.07.003>
- Iolov A, Ditlevsen S, Longti A (2014) Fokker-Planck and Fortet equation-based parameter estimation for a leaky integrate-and-fire model with sinusoidal and stochastic forcing. *J Math Neurosci* 4:4. <https://doi.org/10.1186/2190-8567-4-4>
- Izhikevich EM (2007) *Dynamical systems in neuroscience: the geometry of excitability and bursting*. MIT Press, Boston, MA
- Kang XS, Liang XM, Lü HP (2013) Enhanced response to subthreshold signals by phase noise in a Hodgkin-Huxley neuron. *Chin Phys Lett* 30(1):018701–018704
- Kang Y, Liu R, Mao X (2020) Aperiodic stochastic resonance in neural information processing with Gaussian colored noise. *Cogn Neurodyn* 15(3):517–532
- Koenker R, Mizera I Density estimation by total variation regularization, in: V. Nair (Ed.), *Advances in Statistical Modeling and Inference*, Vol. 3 of Series in Biostatistics, World Scientific, 207, pp. 613–633. https://doi.org/10.1142/9789812708298_0030
- Kraichnan RH (1987) Eddy viscosity and diffusivity: Exact formulas and approximations. *Complex Syst* 1:805–820
- Liang X, Dhamala M, Zhao L, Liu Z (2010) Phase-disorder-induced double resonance of neuronal activity. *Phys Rev E* 82(1):010902–010905
- Liang X, Zhao L, Liu Z (2010) Phase-noise-induced resonance in a single neuronal system. *Phys Rev E* 84(3):031916–031920
- Mackay D (2003) *Information Theory, Inference and Learning Algorithms*, 6th edn. Cambridge University Press, Cambridge
- Maltba TE, Gremaud PA, Tartakovsky DM (2018) Nonlocal PDF methods for Langevin equations with colored noise. *J Comput Phys* 367:87–101
- Morris C, Lecar H (1981) Voltage oscillations in the barnacle giant muscle fiber. *Biophys J* 35(1):193–213
- Neckel T, Rupp F (2013) *Random differential equations in scientific computing*. De Gruyter, Berlin
- Ning LJ, Liu P (2016) The effect of sine-wiener noises on transition in a genotype selection model with time delays. *Eur Phys J B* 89(9):201
- Øksendal B (2003) *Stochastic differential equations*, 6th edn. Springer, Berlin Heidelberg
- Rieke F (1997) *Spikes: exploring the neural code*. Computational neuroscience. MIT Press
- Risken H, Frank T (1996) *The Fokker-Planck equation: methods of solution and applications*. Springer Series in Synergetics, Springer, New York
- Rodgers A, Dektor A, Venturi D (2021) Adaptive integration of nonlinear evolution equations on tensor manifolds. [arXiv:2008.00155](https://arxiv.org/abs/2008.00155)
- Rosenbaum R (2001) A diffusion approximation and numerical methods for adaptive neuron models with stochastic inputs. *Front Comput Neurosci* 38–40:627–632
- Rutjens R, Jacobs G, Tartakovsky DM (2020) Method of distributions for systems with stochastic forcing. *Int J Uncert Quant*. <https://doi.org/10.1615/Int.J.UncertaintyQuantification.2020031940>
- Savchuk OY, Hart JD (2017) Fully robust one-sided cross-validation for regression functions. *Comput Stat* 32(3):1003–1025
- Savchuk O (2020) One-sided cross-validation for nonsmooth density functions, [arXiv:1703.05157](https://arxiv.org/abs/1703.05157)
- Schaeffer H (2017) Learning partial differential equations via data discovery and sparse optimization. *Proc Roy Soc*. <https://doi.org/10.1098/rspa.2016.0446>
- Strand JL (1970) Random ordinary differential equations. *J Diff Equat* 7:538–553
- Tartakovsky DM, Gremaud PA (2015) Method of distributions for uncertainty quantification, In: e. a. R. Ghanem (Ed.), *Handbook of Uncertainty Quantification*, Springer, pp 1–22
- Tuckwell HC, Wan F (2005) Time to first spike in stochastic Hodgkin-Huxley systems. *Physica A: Stat Mech App* 351(2):427–438
- van Kampen NG (1976) *Stochastic differential equations*. *Phys Rep* 24:171–228
- Vejnarová S (1999) The multiinformation function as a tool for measuring stochastic dependence. In: Jordan M (ed) *Learning in graphical models*. MIT Press, pp 261–296
- Venturi D, Sapsis TP, Cho H, Karniadakis GE (2012) A computable evolution equation for the joint response-excitation probability density function of stochastic dynamical systems. *Proc R Soc A* 468(2139):759–783
- Venturi D, Tartakovsky DM, Tartakovsky AM, Karniadakis GE (2013) Exact PDF equations and closure approximations for advective-reactive transport. *J Comput Phys* 243:323–343. <https://doi.org/10.1016/j.jcp.2013.03.001>
- Wang P, Tartakovsky AM, Tartakovsky DM (2013) Probability density function method for Langevin equations with colored noise. *Phys Rev Lett* 110(14):140602
- Watanabe S (1960) Information theoretical analysis of multivariate correlation. *IBM J Res Develop* 4:66–82
- White JA, Rubinstein JT, Kay AR (2000) Channel noise in neurons. *Trends Neurosci* 23(3):131–137
- Yamakou ME, Tran TD, Duc LH, Jost J (2019) The stochastic Fitzhugh-Nagumo neuron model in the excitable regime embeds a leaky integrate-and-fire model. *J Math Biol* 79:509–532
- Yang H, Ning LJ (2017) Phase transitions induced by time-delay and different noises. *Nonlinear Dyn* 88(4):2427–2433
- Yao Y, Ma J (2018) Weak periodic signal detection by sine-Wiener-noise-induced resonance in the FitzHugh-Nagumo neuron. *Cogn Neurodyn* 12:343–349
- Ye M, Neuman SP, Guadagnini A, Tartakovsky DM (2004) Nonlocal and localized analyses of conditional mean transient flow in bounded, randomly heterogeneous porous media. *Water Resour Res* 40:W05104. <https://doi.org/10.1029/2003WR002099>

Zhou X, Xu Y, Wang G, Jia Y (2020) Ionic channel blockage in stochastic Hodgkin-Huxley neuronal model driven by multiple oscillatory signals. *Cogn Neurodyn* 14(4):569–578

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.