

## gPC-based robustness analysis of neural systems through probabilistic recurrence metrics

Uros Sutulovic <sup>a</sup>, Daniele Proverbio <sup>a</sup>, Rami Katz <sup>b</sup>, Giulia Giordano <sup>a</sup>\*

<sup>a</sup> Department of Industrial Engineering, University of Trento, via Sommarive, 9, Trento, 38123, Italy

<sup>b</sup> School of Electrical and Computer Engineering, Tel Aviv University, Chaim Levanon St, 55, Tel Aviv-Yafo, 6997801, Israel

### ARTICLE INFO

Dataset link: <https://github.com/Uros-S/gPC-n-euro-prob-robustness>

#### Keywords:

Dynamical systems  
Dynamical regime  
Regime preservation  
Probabilistic uncertainty  
Probabilistic robustness  
Generalised polynomial chaos  
Surrogate models  
Recurrence plots  
Blob count  
Neuroscience  
Neuronal models  
Neural signals

### ABSTRACT

Neuronal systems often preserve their characteristic functions and signalling patterns, also referred to as *regimes*, despite parametric uncertainties and variations. For neural models having uncertain parameters with a known probability distribution, probabilistic robustness analysis (PRA) allows us to understand and quantify under which uncertainty conditions a regime is preserved in expectation. We introduce a new computational framework for the efficient and systematic PRA of dynamical systems in neuroscience and we show its efficacy in analysing well-known neural models that exhibit multiple dynamical regimes: the Hindmarsh–Rose model for single neurons and the Jansen–Rit model for cortical columns. Given a model subject to parametric uncertainty, we employ generalised polynomial chaos to derive mean neural activity signals, which are then used to assess the amount of parametric uncertainty that the system can withstand while preserving the current regime, thereby quantifying the regime's robustness to such uncertainty. To assess persistence of regimes, we propose new metrics, which we apply to recurrence plots obtained from the mean neural activity signals. The overall result is a novel, general computational methodology that combines recurrence plot analysis and systematic persistence analysis to assess how much the uncertain model parameters can vary, with respect to their nominal value, while preserving the nominal regimes in expectation. We summarise the PRA results through probabilistic regime preservation (PRP) plots, which capture the effect of parametric uncertainties on the robustness of dynamical regimes in the considered models.

### 1. Introduction and motivation

In systems biology, robustness refers to the ability of a system to preserve a qualitative property despite uncertainties, variability or environmental fluctuations [1–4]; also the study of robustness in complex networks [5] offers insight into living systems [6] and neuronal networks [7]. In particular, systems in neuroscience can robustly preserve complex signalling patterns, often called *regimes*, both at the single neuron level and at the neural population level [8–12], despite being subject to environmental disturbances and intrinsic uncertainties [13,14]. Understanding the mechanisms that guarantee such robustness, as well as quantifying the uncertainty levels that are compatible with preservation of a given life-sustaining signalling pattern, enables a better interpretation of recorded signals and a deeper insight into neural dynamics and neural diseases. The study of dynamical regimes and of their robustness in neural models also supports the robust design of artificial systems inspired by dynamics in neuroscience [15], as well as future biomedical applications and medical interventions, e.g., to monitor depth of anaesthesia [16], treat Parkinson's disease [17] and epilepsy [18], and contrast epileptic seizures [19].

\* Corresponding author.

E-mail address: [giulia.giordano@unitn.it](mailto:giulia.giordano@unitn.it) (G. Giordano).

<https://doi.org/10.1016/j.chaos.2026.117949>

Received 7 January 2026; Received in revised form 12 January 2026; Accepted 18 January 2026

Available online 24 January 2026

0960-0779/© 2026 The Authors. Published by Elsevier Ltd. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

The analysis of interpretable dynamical models of neural activity is crucial to understand neural behaviours [20]. To assess how varying key parameters, or combinations thereof, affects the behaviour of neural systems and the resulting (number and type of) observed regimes (such as equilibria), existing works usually employ deterministic methods, such as sensitivity analysis, parameter sweeping or bifurcation analysis [21,22], where parameter values are assumed to be constant and deterministic. These methods to assess the number and properties of dynamical regimes in a parameter space, and the regime changes associated with parameter variations, do not account for the possibility that biological parameters are affected by probabilistic uncertainties. Also, given experimental data, traditional modelling typically employs deterministic methods to identify the single best-fit value of the system parameters that matches the observed biological phenomenon and then uses the resulting model for prediction; the usefulness and reliability of the obtained predictions critically depend on the accuracy, precision and completeness of measurements [23]. However, parametric uncertainty and variability strongly affect real-life neuronal systems, and thus the ensuing measurements, as observed in experimental settings [24]. Probabilistic parametric uncertainty causes the system parameters to take values according to a probability distribution, within feasible bounds including a nominal value, thus yielding different system realisations (such as those observed in distinct neurons within a population) that may exhibit the same qualitative physical behaviour despite having different system parameters. This exactly matches real-life scenarios: different cells, neurons and structures with the same function do not typically share the same parameter values, since these are associated with physiological processes, such as ion fluxes through membrane channels, that are subject to variability and thus induce heterogeneity, across cells or over time [24,25]. Uncertainty and variability, therefore, are intrinsic and cannot be neglected in the study of neuronal systems.

Mathematically, uncertainty can be characterised either in a *deterministic* or in a *probabilistic* setting. Studying the effect of *deterministic* parametric uncertainties is the goal of classic robustness analysis, which has a long tradition in control theory [26]. Robustness analysis is agnostic about the exact parameter values realised in practice, and provides a “yes” or “no” answer to the question whether a property of interest (in our case, a regime) is preserved robustly, *i.e.*, for all possible parameter realisations within a known bounding set in the parameter space. The resulting sufficient conditions for the desired system performance are often overly conservative [26, Chapter 14] and have an “all-or-nothing” nature: if, even for a single parameter choice within the bounding set, the property does not hold, then robustness cannot be guaranteed. In nature, however, certain proportions of defective cells and processes may be present while proper biological functions are maintained nonetheless [27]. Hence, preservation of a life-sustaining behaviour “on average” (*e.g.*, through a critical mass of properly functioning cells) is what truly matters. This observation fosters the need to investigate *probabilistic* parameter variations in biological systems, and their effects on the preservation of biological functions and regimes.

In this work, to study the robustness of neurological systems with respect to parametric uncertainty, we resort to *probabilistic* robustness analysis (PRA) [28]. Differently from deterministic methods such as sensitivity and bifurcation analysis, and from classic parametric robustness analysis, PRA allows us to consider probabilistic parametric uncertainty in neurological models and more accurately analyse the effects of randomness in parameter realisations on the resulting (stochastic) behaviour of neural systems. In fact, PRA of a system whose uncertain parameters can vary, within a known bounding set in the parameter space, according to a known probability distribution (which encodes likelihood belief, estimation error, or population variability) has several advantageous properties.

(i) PRA identifies uncertainty regions in the parameter space where the same qualitative system behaviour is preserved [29–32] in probability. PRA can not only assess whether a system preserves a regime of interest in expectation in the presence of *given* parametric uncertainties affecting its nominal parameters, but also *quantify how much* parametric uncertainty affecting its nominal parameters a system can withstand, while still preserving a regime of interest in expectation. (ii) PRA helps interpret the emergent *average* effects of probabilistic parametric uncertainty: it allows to quantify, predict and analyse the *mean* neural activity of ensembles of heterogeneous neurons, which is considered to be a more robust and reliable proxy for neural functions than the activity of single neurons, individually more fragile and prone to malfunctioning [33]. Importantly, the signal associated with nominal deterministic values of the parameters may be completely different from the mean signal realisation obtained when the parameters are subject to parametric uncertainty within a parameter bounding set that contains the nominal parameter values (see Supplementary Figure S4); hence, for instance, regime identification is much more difficult, compared to the deterministic setting. If all parameter realisations in the bounding set yield the same qualitative behaviour (as in deterministic robustness analysis), then the mean realisation in the set is more likely, but is not guaranteed, to preserve the same qualitative behaviour. These observations make the question of mean regime preservation under probabilistic parametric uncertainty much more intricate than its deterministic robustness counterpart. In particular, studying the *deterministic* realisation corresponding to the *average* of multiple values of the system parameters cannot provide any meaningful information on the *mean* system realisation. Indeed, for systems in neuroscience it is well known that the signal associated with the mean of multiple parameter values may not be representative of the regime associated with each of the original parameter values [34] (see, *e.g.*, Supplementary Figure S3). Systematic PRA tackles this complexity. (iii) When comparing models to data, *in-vivo* recording methods such as electro-encephalograms (EEG) or local field potentials (LFP) measure the mean signals over batches of neurons [35,36]. Therefore, PRA is a promising tool to understand whether regimes obtained from single-neuron models can robustly predict observed ensemble recordings, and to improve the synergy between mathematical models and real-life data through more reliable parameter fitting. (iv) The ability to detect models that reproduce observed neurological phenomena only for specific parameter choices and cannot accommodate even reasonably small parameter variations while preserving the intended neurological behaviour, as opposed to models that display robust neurologically-meaningful regimes in the face of natural parametric uncertainty, makes PRA precious to falsify models and guide model selection, also in combination with suitably designed experiments [29].

In this paper, we develop a novel methodology for the efficient and systematic PRA of uncertain neuronal systems and we demonstrate its efficacy on well-established dynamical models from mathematical neuroscience: the Hindmarsh–Rose (HR) model [37] and the Jansen–Rit (JR) model [38]. We consider a class of models of the form  $\dot{\mathbf{x}} = f(\mathbf{x}, Z)$  subject to probabilistic parametric uncertainty, where the parameters in  $Z$  are independent random variables characterised by a known uniform probability distribution, supported on a bounded set that contains the nominal parameter values. Then, the state  $\mathbf{x}(t; Z)$  is a stochastic process and each realisation of the system trajectory depends on a realisation of the parameters in  $Z$ . In this probabilistic framework, we aim to identify regions in the parameter space where the regime associated with the deterministic nominal model is preserved in expectation despite parametric uncertainties, so as to distinguish between regimes and quantify the largest uncertainty level compatible with probabilistic robustness of a regime of interest. To assess probabilistic robustness, we focus on time-averages (means) of output signals of interest (obtained from the imposed parameter distributions); to efficiently compute mean neural activities under parametric uncertainty, we resort to the generalised polynomial chaos (gPC) method [39], which has two major advantages: first, it effectively incorporates parametric uncertainty into the analysed models and provides explicit analytic expressions that directly capture the uncertainty effects; second, it provides a favourable trade-off between accuracy and computational costs [40], often being much more efficient than popular Monte Carlo methods, a commonly used alternative for model-based PRA. Then, we propose a computational methodology to extract patterns from the obtained mean time-series through the generation of recurrence plots [41]. The identification of distinct geometric patterns within recurrence plots enables to distinguish between regimes; we show that our proposed approach, based on *blob counts*, successfully copes with the presence of parametric uncertainty (while other metrics often used for regime identification in deterministic models, such as inter-spike intervals, duty cycles or spike counts, fail to identify regimes and to infer reliable information about their preservation). Finally, our methodology produces *probabilistic regime preservation* (PRP) plots that contain essential information on the dynamical regimes induced by the system under study, and their robustness to parametric uncertainties, in a parameter space of interest. Differently from classical deterministic methods, such as bifurcation plots, the obtained PRP plots offer a probabilistic perspective on the mean behaviour of the system under investigation, and on the preservation of its regimes when the system parameters are affected by probabilistic uncertainty.

We present in detail our proposed computational methodology for PRA of neural systems in Section 2, after providing some necessary background on the considered models and on the gPC technique. Then, in Section 3, we apply our new approach to the HR and the JR models, and discuss the significance of the obtained PRP results for their robustness to parametric uncertainty. A concluding discussion is given in Section 4. Additional information and simulations are provided in the Supplementary Material.

## 2. Methodology: assessing regime preservation under parametric uncertainty

This section describes our novel methodology for the PRA of systems in neuroscience subject to parametric uncertainties, leading to the calculation of PRP plots, as shown in Fig. 1.

### 2.1. Class of considered systems subject to parametric uncertainties

We consider models from mathematical neuroscience that are described by an ordinary differential equation (ODE) system of the form  $\dot{\mathbf{x}} = f(\mathbf{x}, Z)$ , where  $\mathbf{x} \in \mathbb{R}^n$  is the state vector and the uncertain system parameters of interest are stacked in a parametric uncertainty vector  $Z \in \mathbb{R}^d$ . We assume that the components of  $Z$  are independent random variables, characterised by a known probability distribution supported on a bounded set representing the parameter space of interest for the considered model. Many different distributions of  $Z$  may be handled by our proposed approach [39]. Since the uniform distribution has been observed to be the worst-case distribution for PRA under suitable assumptions [42], and we can theoretically transform any continuous random variable into a uniformly distributed one through a quantile transformation [43], in the absence of additional information on  $Z$ , we assume that each component  $Z_i, i = 1, \dots, d$ , is *uniformly* distributed over an uncertainty interval containing the nominal value  $Z_i^*$ . Precisely, we therefore have  $Z_i \sim \mathcal{U}([Z_{i,\min}, Z_{i,\max}])$ , with  $Z_{i,\min} \leq Z_i^* \leq Z_{i,\max}, i = 1, \dots, d$ . The resulting ODE model subject to parametric uncertainty is

$$\dot{\mathbf{x}} = f(\mathbf{x}, Z), \quad Z \sim \mathcal{U}(\mathcal{A}), \quad \mathcal{A} = [Z_{1,\min}, Z_{1,\max}] \times \dots \times [Z_{d,\min}, Z_{d,\max}] \subset \mathbb{R}^d, \tag{1}$$

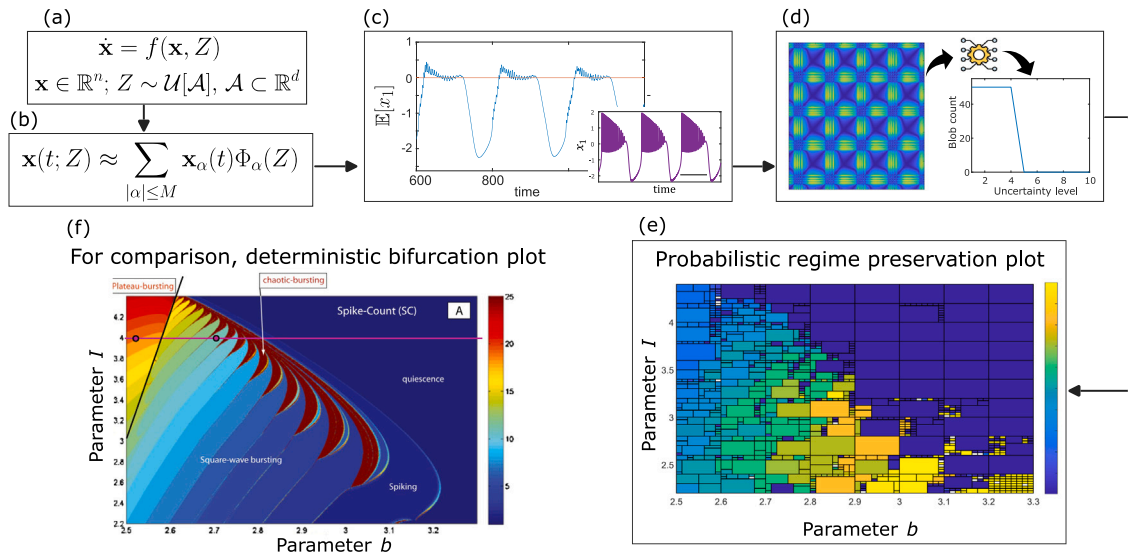
where  $f : \mathbb{R}^n \times \mathbb{R}^d \rightarrow \mathbb{R}^n$ , the hyper-rectangle  $\mathcal{A}$  is the uncertainty bounding set and  $\mathbf{x}(t; Z)$  is a stochastic process (Fig. 1a); in this probabilistic setting, different system evolutions are possible depending on the specific realisation of the parameters in  $Z$ . In (1), stochasticity arises from probabilistic uncertainty in the system parameters, and not from a stochastic process driving the dynamics as it happens for instance in stochastic differential equations [44].

Our approach can be applied to any system subject to parametric uncertainty, whose parameter space is not too high-dimensional, since the only bottleneck relates to computational cost.

To demonstrate our proposed methodology, we focus on two celebrated models from neuroscience: the Hindmarsh–Rose (HR) model for single-neuron dynamics [37] and the neural mass mean-field Jansen–Rit model for a single cortical column [38].

The HR model [37] is widely used to reproduce the spiking and bursting dynamics of single neurons observed in experimental data [45]; however, its robustness to probabilistic parametric uncertainty has not been studied so far. In the adimensionalised ODEs of the HR model

$$\begin{cases} \dot{x}_1(t) = x_2(t) - ax_1^3(t) + bx_1^2(t) - x_3(t) + I, \\ \dot{x}_2(t) = c - dx_1^2(t) - x_2(t), \\ \dot{x}_3(t) = r[s(x_1(t) - x_R) - x_3(t)], \end{cases} \tag{2}$$



**Fig. 1.** Workflow of our methodology for probabilistic robustness analysis (PRA) described in Section 2. (a) Consider a dynamical system of interest, where the evolution of the state vector  $x \in \mathbb{R}^n$  depends on a vector  $Z$  of  $d$  uncertain parameters, independent and distributed according to uniform distributions within known bounding intervals. (b) Derive a surrogate model of the system with the gPC method. (c) Based on the surrogate model, efficiently compute the mean system output (blue) with respect to the parametric uncertainty  $Z$ ; the corresponding deterministic output of the nominal, uncertainty-free system is shown (purple) for comparison. (d) To unravel key geometric patterns within the time series associated with the mean system output, construct the corresponding recurrence plot (left). Then, automatically extract persistent features (blob count) of recurrence plots corresponding to increasing uncertainty levels, so as to assess regime preservation: the regime is preserved for all the uncertainty levels that achieve a similar blob count, while it is no longer preserved once the blob count drops (right). (e) Capture the overall information about dynamical regimes and their robustness to parametric uncertainty in a probabilistic regime preservation (PRP) plot, where the colour encodes blob counts and the length of each side of the rectangles quantifies how much uncertainty the system can withstand in the corresponding parameter value while maintaining the same regime in expectation. (f) For comparison, in the same parameter space as the PRP plot, we present a bifurcation plot (image courtesy of Springer, see [21] for additional details) that identifies the regimes resulting from different deterministic values of the parameters.

variable  $x_1(t)$  represents the neuron membrane potential,  $x_2(t)$  a fast recovery current, and  $x_3(t)$  a slow adaptation current that enables bursting (*i.e.*, trains of multiple membrane potential spikes).  $I$  is an injected current,  $x_R$  is the reference resting membrane potential,  $r$  is the timescale ratio between fast and slow currents, while parameters  $a, b, c, d, s$  are related to the system nullclines [37]. We consider  $Z_{HR} = [b, I]^T$  as the parameters of interest, and set the other parameters to the standard values in [37]:  $a = 1, c = 1, d = 5, s = 4, x_R = -8/5, r = 0.01$ . In a *deterministic* setting, for different values of the system parameters, the state trajectories  $x(t) = [x_1(t), x_2(t), x_3(t)]^T, t \geq 0$ , of this relatively simple model can exhibit multiple qualitatively distinct behaviours, which we call *dynamical regimes*. We focus on the membrane potential,  $y(t) = x_1(t)$ , as an output signal of interest; examples of realisations of  $x_1(t)$  for different parameter values that pertain to different dynamical regimes, ranging from quiescence to bursting, are shown in Fig. 2 (top).

The JR model [38] is a neural mass mean-field model that describes, rather than the dynamics of individual neurons, the generation of EEG signals by a single cortical column. In particular, the cortical column is represented through the interaction of the mean-field activity of three neuronal populations: pyramidal neurons, excitatory interneurons and inhibitory interneurons. In the corresponding ODE system

$$\begin{cases} \dot{x}_1 = x_4, & \dot{x}_4 = AaS(x_2 - x_3) - 2ax_4 - a^2x_1, \\ \dot{x}_2 = x_5, & \dot{x}_5 = Aa[p + C_2S(C_1x_1)] - 2ax_5 - a^2x_2, \\ \dot{x}_3 = x_6, & \dot{x}_6 = BbC_4S(C_3x_1) - 2bx_6 - b^2x_3, \end{cases} \quad (3)$$

the states  $x_i(t), i = 1, 2, 3$ , are neural potentials ( $x_1$  is the mean excitatory postsynaptic potential,  $x_2$  is the mean membrane potential of pyramidal cells,  $x_3$  is the mean inhibitory postsynaptic potential), while  $S(V) = v_M / (1 + e^{r(V_0 - V)})$  is a sigmoidal function, where  $v_M$  is the maximal firing rate of the family of neurons,  $V_0$  is the excitability threshold of the population,  $r$  is the firing threshold variance (*i.e.*, the slope of the sigmoid at  $V_0$ ). Parameters  $A$  and  $B$  are the maximum amplitude of excitatory and inhibitory postsynaptic potentials, respectively, while  $a$  and  $b$  are lumped parameters related to time constants in the dendritic network for excitatory and inhibitory synapses, respectively. All connectivity constants  $C_i$ , describing the strength of the links between pyramidal neurons and inhibitory and excitatory interneurons, are proportional to a single parameter  $C$ :  $C_1 = C, C_2 = 0.8C, C_3 = 0.25C, C_4 = 0.25C$ , with  $C = 135$  [38]. We investigate  $Z_{JR} = [A, B]$  as the parameters of interest, for varying values of  $p$  and  $C$ , while all other parameters

are fixed as in [38]:  $a = 100, b = 50, V_0 = 6, v_M = 5, r = 0.56$ . In a deterministic setting, the EEG-like output signal  $y(t) = x_2(t) - x_3(t)$  exhibits different qualitative dynamical behaviours (regimes) as shown in Supplementary Figure S2.

We consider systems (2) and (3) subject to probabilistic parametric uncertainties as in (1). In particular, the HR model can be cast in the form (1) with  $\mathbf{x} = [x_1(t), x_2(t), x_3(t)]^T, Z = Z_{HR} = [b, I]^T, f(\mathbf{x}, Z_{HR})$  given by the right-hand side of (2) and  $\mathcal{A}$  being a rectangle that includes the nominal parameter values  $Z_{HR}^* = [b^*, I^*]^T$ . Similarly, for the JR model,  $\mathbf{x} = [x_1(t), x_2(t), x_3(t), x_4(t), x_5(t), x_6(t)]^T, Z = Z_{JR} = [A, B]^T, f(\mathbf{x}, Z_{JR})$  is given by the right-hand side of (3) and  $\mathcal{A}$  is a rectangle that includes the nominal parameter values  $Z_{JR}^* = [A^*, B^*]^T$ .

To illustrate the next steps of our methodology, in this section we focus on the HR model (2); the results for the JR model (3) are presented in Section 3.

We consider the HR parameter space  $[b_{\min}, b_{\max}] \times [I_{\min}, I_{\max}] = [2.5, 3.3] \times [2.2, 4.4]$ . Fig. 1f shows the corresponding bifurcation diagram, from [21], in the deterministic setting: regions with different colours correspond to deterministic parameter values that result in regimes with different characteristics (exemplified in Fig. 2). In our numerical experiments, we choose nominal parameter values  $[b^*, I^*]^T$  that lie within the interior of a region associated with one of the deterministic nominal regimes in Fig. 1f. Probabilistic parametric uncertainty affecting the nominal parameter values may lead to loss of regime identifiability from the mean output trajectory realisation (see Fig. 1c and Supplementary Figure S4 for examples). Our goal is thus to develop a computational framework that allows us to detect regimes in the presence of parametric uncertainty and identify the highest possible uncertainty level (corresponding to the largest possible neighbourhood of the nominal parameters, given predefined specifications) where the nominal regime, exhibited by the system output under the deterministic nominal parameters, is preserved in expectation. To this aim, we design dedicated metrics and apply them to the mean output signal, so as to relate it to the nominal regime.

### 2.2. Generalised polynomial chaos (gPC) surrogate models

To compute statistics of the stochastic process  $\{\mathbf{x}(t; Z)\}_{t \geq 0}$  described by (1), we employ a gPC approximation [39] to obtain a surrogate model of (1), based on polynomial regression. The gPC surrogate model is simpler to integrate and retains the probabilistic information needed to efficiently calculate the mean process  $\{\mathbb{E}[\mathbf{x}(t; Z)]\}_{t \geq 0} =: \{\mathbb{E}[\mathbf{x}](t)\}_{t \geq 0}$ , where  $\mathbb{E}$  denotes the expectation with respect to  $Z \sim \mathcal{U}(\mathcal{A})$ . This gPC-based approach has been recently shown to provide surrogate models characterised by a favourable trade-off between computational effort and accuracy, for relevant systems in neuroscience at different spatial scales [40].

To introduce the gPC methodology, consider a general ODE system (1) subject to parametric uncertainty. The state trajectory  $\{\mathbf{x}(t; Z)\}_{t \geq 0}$  is a stochastic process that depends on the parametric uncertainty  $Z \in \mathbb{R}^d$ , and we wish to efficiently compute the mean state trajectory  $\mathbb{E}[\mathbf{x}] : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}^n$ , where  $\mathbb{E}[\mathbf{x}](t) := \frac{1}{\text{size}(\mathcal{A})} \int_{\mathcal{A}} \mathbf{x}(t; \xi) d\xi$  and  $\text{size}(\mathcal{A}) = \prod_{i=1}^d (Z_{i,\max} - Z_{i,\min})$ .

We assume that the stochastic process  $\{\mathbf{x}(t; Z)\}_{t \geq 0}$  has finite variance, which is not restrictive for mathematical models in neuroscience, since any meaningful model would take into account biophysical constraints, thus producing bounded state trajectories. Specifically, we assume that for each  $t \geq 0$  (outside perhaps of a zero measure set), every component of  $\mathbf{x}(t; \cdot)$ , treated as a function of  $Z$ , belongs to the space  $L^2_{\mathcal{U}}(\mathcal{A}) = \{\psi : \mathbb{E}[\psi^2] := \frac{1}{\text{size}(\mathcal{A})} \int_{\mathcal{A}} \psi^2(\xi) d\xi < +\infty\}$ . This is guaranteed if the function  $\mathbf{x}(t; \cdot)$  is measurable and bounded. Since the components of  $Z$  are independent, a natural basis for  $L^2_{\mathcal{U}}(\mathcal{A})$  is the set of multivariable polynomials  $\{\Phi_{\alpha}\}_{\alpha}$  that satisfy the orthogonality relation  $\mathbb{E}[\Phi_{\alpha} \Phi_{\beta}] = \frac{1}{\text{size}(\mathcal{A})} \int_{\mathcal{A}} \Phi_{\alpha}(\xi) \Phi_{\beta}(\xi) d\xi = \delta_{\alpha\beta}$ , where  $\delta_{\alpha\beta}$  is the  $d$ -variate Kronecker delta and  $\alpha = (\alpha_1, \dots, \alpha_d) \in \mathbb{N}_0^d$  is a multi-index. The basis of orthogonal polynomials  $\{\Phi_{\alpha}\}_{\alpha}$  can be computed via a Gram–Schmidt orthogonalisation process [39]. As proven in [46], the following series representation holds in  $L^2_{\mathcal{U}}(\mathcal{A})$ :

$$\mathbf{x}(t; Z) = \sum_{\alpha \in \mathbb{N}_0^d} \mathbf{x}_{\alpha}(t) \Phi_{\alpha}(Z), \tag{4}$$

where the choice of the basis  $\{\Phi_{\alpha}\}_{\alpha}$  depends on the probability distribution of the random variables in  $Z$ , according to the Wiener–Askey table [39,47] that associates probability distributions with classes of orthogonal polynomials. In our case, since the random variables in  $Z$  are uniformly distributed,  $\{\Phi_{\alpha}\}_{\alpha}$  is composed of Legendre polynomials. Importantly, the series representation (4) neatly decouples the effect of the uncertain parameters  $Z$ , confined to the basis of orthogonal polynomials  $\{\Phi_{\alpha}(Z)\}_{\alpha}$ , and the deterministic time dependence, contained in the spectral coefficients  $\{\mathbf{x}_{\alpha}(t)\}_{\alpha}$ . This greatly simplifies the computation of statistics of the random process  $\mathbf{x}(t; Z)$ ; for instance, mean and variance are given by

$$\begin{cases} \mathbb{E}[\mathbf{x}](t) = \mathbf{x}_0(t), \\ \sigma_{\mathbf{x}}^2(t) = \sum_{|\alpha| \neq 0} \mathbf{x}_{\alpha}^2(t). \end{cases} \tag{5}$$

To obtain a tractable surrogate model for practical implementation (Fig. 1b), we truncate the series (4) to an expansion order  $M$ , namely, we only retain indices with  $|\alpha| := \sum_{i=1}^d \alpha_i \leq M$ , thus approximating  $\mathbf{x}(t; Z)$  with  $\mathbf{x}_M(t; Z) = \sum_{|\alpha| \leq M} \mathbf{x}_{\alpha}(t) \Phi_{\alpha}(Z)$ . Therefore, construction of the surrogate model only requires finding the spectral coefficients  $\{\mathbf{x}_{\alpha}\}_{|\alpha| \leq M}$ . For this, we adopt the non-intrusive collocation approach, which consists of computing  $S_C \geq \binom{M+d}{d}$  system trajectory realisations and then estimating the spectral coefficients indirectly, e.g., via a least-squares approximation [40,48]. Since theoretical convergence guarantees are available for static maps of random variables [49], but not for solutions to general ODE systems, it is important to numerically check the suitability of the gPC-based approximation, by computing the coefficients  $\{\mathbf{x}_{\alpha}(t)\}_{|\alpha| \leq M}$  with increasing values of  $M$ , until the decimal expansion of a fixed initial set of coefficients  $\{\mathbf{x}_{\alpha}(t)\}_{|\alpha| \leq M_0}$ , with  $M_0 < M$ , stabilises to a prescribed order of decimal digits. For computational purposes, we implement gPC collocation via the PoCET toolbox in MATLAB [50].

### 2.3. Mean state trajectories and output signals

Given the parametric uncertainty  $Z$ , the gPC approach allows us to efficiently compute the mean system trajectory  $\mathbb{E}[\mathbf{x}](t)$  as in (5). The mean output trajectory  $\mathbb{E}[y](t)$ , where  $y \in \mathbb{R}$  is a linear combination of the state variables, can then be obtained from  $\mathbb{E}[\mathbf{x}](t)$ . For the HR model, Fig. 1c shows the time evolution of the deterministic output  $x_1(t; Z^*)$  (inset, purple) with nominal parameters  $Z^* = [b^*, I^*]^T = [2.44, 4.2]^T$  associated with the plateau bursting regime, along with the mean system output  $\mathbb{E}[x_1](t)$  (blue) obtained when  $I = I^* = 4.2$  and  $b \sim \mathcal{U}([2.4, 2.48])$ .

Sampling the mean system output yields a time-series  $\{\mathbb{E}[y](t_k)\}_k$ , formed by a collection of average values at time instants  $\{t_k\}_k = \{k \cdot \Delta t\}_k$ ,  $k \in \mathbb{N}$ , where  $\Delta t > 0$  is a fixed inter-sampling time. Then, we need to design criteria to determine whether, in the presence of the uncertainty  $Z$ , the sampled mean output signal  $\{\mathbb{E}[y](t_k)\}_k$  preserves the qualitative features that characterise the nominal regime. Assessing regime preservation in the presence of probabilistic parametric uncertainty is particularly challenging: we cannot resort to the standard metrics employed in the deterministic case for regime identification, as for mean signals they either cannot be computed or lose their interpretability (see Supplementary Section S2). In the deterministic case, i.e., for fixed parameter values, standard simple metrics such as inter-spike interval, duty cycle or spike count, often employed to extract quantitative information from neural signals [20,21,51], are easy to compute. However, with probabilistic uncertainty, the mean output signal is obtained by averaging different realisations of the output signal, whose spikes and bursts occur at different time instants or intervals. Hence, standard metrics are not properly defined for mean signals, where it is not even clear what constitutes a spike or a burst (see Supplementary Figure S5). Thresholding procedures further hamper the generic use of such standard metrics [52]. To distinguish between different regimes, an alternative approach may be to consider the mean output signal as a time-series and compute its power spectral density (or some wavelet transform) to search for distinctive frequency components [53,54] (Supplementary Figure S9). However, through numerical investigations partially reported in Supplementary Section S2, we verified the current inefficacy of these and other existing methods used for regime identification in deterministic models, which become generally inapplicable in the presence of probabilistic parametric uncertainties.

In deterministic realisations of the HR model, the plateau-bursting regime (D) in Fig. 2 is characterised by variable  $x_1$  exhibiting oscillations, initiated after it reaches the value  $x_1 \approx 0$ , which last with gradually damped amplitude for a finite time interval, after which  $x_1 \approx 0$  again. If such a distinctive characteristic is also present in the mean output signal for a given parametric uncertainty (as in Fig. 1c), then we may say that the regime is preserved in a probabilistic sense [40, Section III.C]. Unfortunately, cases in which the mean signal preserves such properties of the corresponding nominal regime are extremely rare. For the HR model, this happens only for plateau-bursting and not for other regimes; for other systems, it may never happen.

Therefore, for regime identification in a probabilistic uncertainty setting, other general techniques or metrics are needed, which we propose next.

### 2.4. Persistence of geometric structures in recurrence patterns

We describe here our novel methodology that integrates recurrence plots and automated blob counting to study preservation of dynamical regimes in expectation under probabilistic parametric uncertainty. We continue to illustrate our methodology on the HR model (2), where we consider as output of interest  $y(t) = x_1(t)$ , the membrane potential, whose distinct qualitative behaviours give rise to different dynamical regimes such as those shown in Fig. 2.

#### 2.4.1. Recurrence analysis

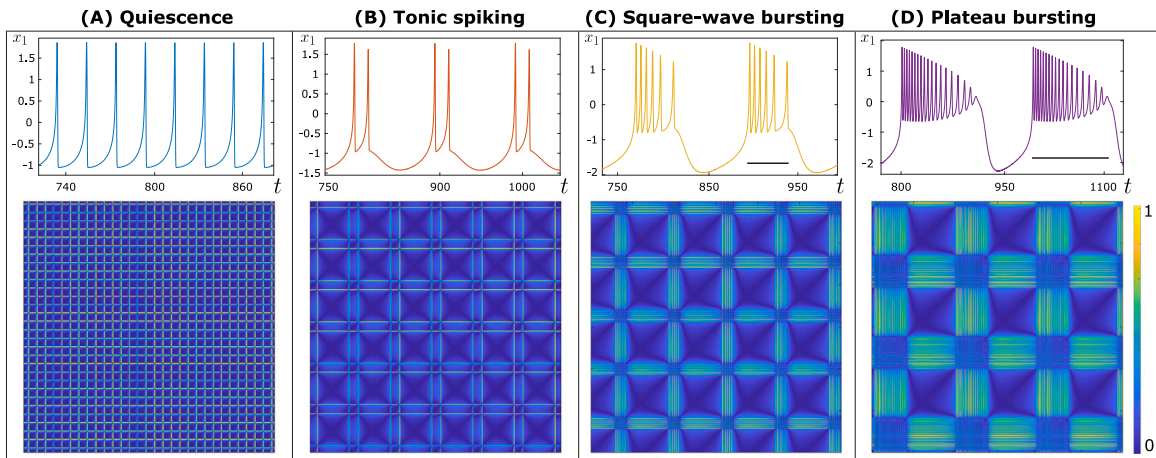
To extract information from mean output signals aimed at assessing regime preservation under parametric uncertainty, we resort to recurrence analysis [55]. An advantage of recurrence plots [41] lies in their ability to effectively visualise characteristic features of signals through constructed geometric patterns, which reveal qualitative signatures of, e.g., periodic and quasi-periodic behaviour, slowly evolving dynamics, stochastic-like fluctuations, or deterministic chaotic regimes [56,57]. To construct the recurrence plot, given  $T$  samples of the mean output  $\{\mathbb{E}[y](t_k)\}_{k=1}^T$  as our discrete-time signal of interest, we compute the distance matrix  $D$  as

$$D_{\ell h} = |\mathbb{E}[y](t_\ell) - \mathbb{E}[y](t_h)|, \quad 1 \leq \ell, h \leq T. \tag{6}$$


The entries of  $D$  in (6) yield an *unthresholded* recurrence plot [56, Subsection 3.2.5], which processes the distance matrix  $D$  as an image with pixels coloured according to the values of its entries. Since the recurrence patterns are preserved under scaling, without loss of generality,  $D$  is normalised by its maximum value, so that all of its entries  $D_{\ell h}$  lie in  $[0, 1]$ . The recurrence plots for the HR model (2), corresponding to distinct *deterministic* realisations of  $y(t) = x_1(t)$ , the membrane potential, in different dynamical regimes, are shown in Fig. 2 (bottom); see also Fig. 1d, left. Recurrence plots allow us to visually distinguish between different dynamical regimes in a clear manner. This is also true for the JR model (3), see Supplementary Figure S2.

To study regime preservation, we analyse how the recurrence plot degrades when we increase the uncertainty level, i.e., the size of the support of the uniform distribution from which the uncertain parameters are randomly chosen. In particular, we consider a sequence of  $N$  recurrence plots obtained for distributions with increasing uncertainty level:  $\{Z \sim \mathcal{U}(\mathcal{A}_i)\}_{i=1}^N$ , where the nominal value  $Z^* \in \mathcal{A}_i$  for all  $i = 1, \dots, N$  and  $\mathcal{A}_{i+1} \supset \mathcal{A}_i$  for  $i = 1, \dots, N - 1$ .

For the HR model (2), we fix  $I = 2.4$  and the nominal  $b^* = 2.7$ , consider uncertain  $b$  with distributions  $\{b \sim \mathcal{U}([b^*, b^* + \Delta b_i])\}_{i=1}^N$  having support of increasing size  $\{\Delta b_i = \frac{i}{N} \Delta b_{\max}\}_{i=1}^N$  up to a maximum size  $\Delta b_{\max}$ , and assess whether the characteristic pattern associated with the nominal square-wave bursting regime persists in the recurrence plots for increasing  $i = 1, \dots, N$ , where the index  $i$  encodes the uncertainty level. Fig. 3 shows the sequence of recurrence plots obtained with  $\Delta b_{\max} = 0.2$  and  $N = 10$ . The



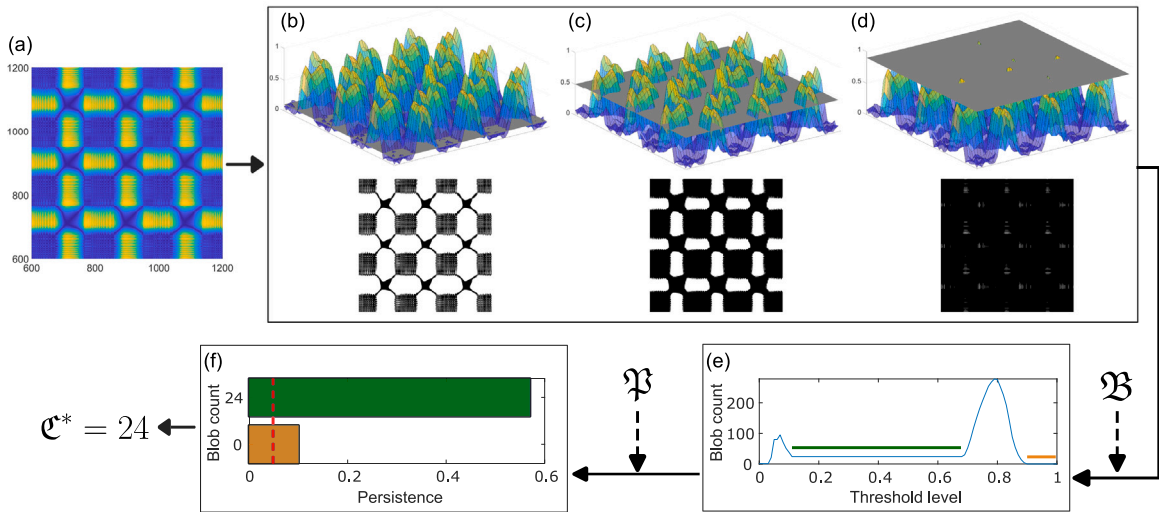
**Fig. 2.** For the HR model (2), deterministic realisations of  $x_1(t)$  (top) and corresponding normalised recurrence plots in the time window [600, 1200] (bottom), with parameters  $b \in [2.5, 3.3]$  and  $I \in [2.2, 4.4]$  that yield different regimes: (A)  $b = 2.8$  and  $I = 4.2$ , quiescence; (B)  $b = 3.1$  and  $I = 2.4$ , tonic spiking; (C)  $b = 2.65$  and  $I = 2.4$ , square-wave bursting; (D)  $b = 2.5$  and  $I = 4.0$ , plateau bursting. In (C) and (D), the black horizontal line identifies a burst of spikes. Note that what we call quiescence may also be called spiking, while what we call tonic spiking may also be called two-bursting [58].



**Fig. 3.** Degradation of recurrence plot patterns associated with the HR square-wave bursting regime in Fig. 2 (C) for increasing uncertainty levels  $i = 1, \dots, N$ , with  $N = 10$ . Here  $I = 2.4$ , while  $b \sim \mathcal{U}([b^*, b^* + \Delta b_i])$  for the  $i$ th uncertainty level, with  $b^* = 2.7$  (in red) and  $\Delta b_i = \frac{i}{N} \Delta b_{\max}$  with  $\Delta b_{\max} = 0.2$  (i.e.,  $\Delta b_i = 0.02i$  for  $i = 1, \dots, 10$ ). The nominal deterministic recurrence plot with  $I = 2.4$  and  $b = b^* = 2.7$  is shown in Supplementary Figure S1.

uncertainty pattern corresponding to the nominal square-wave bursting regime is preserved roughly up to the fourth uncertainty level,  $b \sim \mathcal{U}([2.7, 2.78])$ . Then, the geometric patterns present in the nominal recurrence plot become significantly less prominent, and hence cannot be unequivocally associated with the square-wave bursting regime, at least based on visual inspection; the reason is that the uncertainty region spans a portion of the parameter space associated with trajectories having significantly different behaviours (e.g., with spikes occurring at vastly different time instants) and hence recurrence plots of the mean output signal  $\mathbb{E}[x_1]$  do not retain a distinctive characteristic.

Recurrence analysis allow us to transform the problem of assessing regime preservation into image pattern recognition/retention. This task is often tackled using neural networks [59], which however are not suitable in our case, because training and validation processes do not generalise well from one system to another, require extensive labelling, and are computationally heavy. A possible alternative would be unsupervised classification of images [60], which however requires knowing the number of labels in advance; since we are not tackling the problem of global classification of dynamical regimes in a parameter space, and we are only interested



**Fig. 4.** Automated blob count for (a) the recurrence plot  $D$  of the mean output  $\mathbb{E}[x_1](t)$  for the HR model (2), with  $b = 2.5$  and  $I \sim \mathcal{U}([3.6, 3.8])$ . (b)–(c)–(d) Matrix  $D$  is visualised as a surface and different Boolean versions  $\mathcal{D}$  are obtained by processing with different thresholds: (b)  $\vartheta = 0.1$ ; (c)  $\vartheta = 0.5$ ; (d)  $\vartheta = 0.9$ . (e) We discard blobs smaller than  $\mathfrak{B} = 150$  pixels and show the blob count  $\mathcal{C}$  as a function of  $\vartheta$ . (f) Persistence bar chart; among the blob counts with persistence larger than a user-defined minimum persistence  $\mathfrak{P} = 0.05$  (dashed red line), we select  $\mathcal{C}^* = 24$  as it corresponds to lower thresholds; in fact, as shown in panel (e),  $\mathcal{C} = 24$  is obtained for the thresholds highlighted in green,  $\mathcal{C} = 0$  for those in orange.

in a local classification aimed at assessing local regime preservation under parametric uncertainty, we take a more direct and interpretable approach that is not based on machine learning algorithms, and leave their investigation for future research.

#### 2.4.2. Automated blob counting for recurrence plots

Since recurrence plots exhibit a degradation of the geometric patterns associated with the nominal regime as the uncertainty level (i.e., the support of the uniform distribution) increases, the next step in PRA is to formulate a systematic, computable criterion to determine at which level of parametric uncertainty the geometric patterns in the recurrence plot are considered to be lost.

Let us denote by matrices  $\{D_i\}_{i=1}^N$  as in (6) the recurrence plots associated with the uncertainty intervals  $\{\Delta b_i\}_{i=1}^N$ , respectively. Preliminary attempts for regime identification include computing a chosen matrix-related metric (e.g., maximum value, mean value, standard deviation and root-mean-square of all entries, spectral norm, or Frobenius norm) on each of the  $D_i$ 's and then attempting to associate the loss of the geometric patterns in the recurrence plot with a significant change in the metric (Supplementary Figures S6 and S7). Alternatively, by interpreting the  $D_i$ 's as matrices that describe probability mass functions defined on a common bounded domain, given two such matrices, one may compute the minimum “cost” to turn one into the other through Wasserstein distance, or mutual information quantifying their amount of shared information, or Jaccard index quantifying their degree of similarity (Supplementary Figure S8). For all of these approaches, among the mentioned metrics we found no definitive criterion that allowed us to reliably distinguish between different dynamical regimes (through distinct intervals of the metric values associated with distinct regimes) in the presence of parametric uncertainty, even for relatively small uncertainty levels, and reliably discern the loss of geometric patterns in recurrence plots (through a significant change in the metric) due to a significant increase in the uncertainty level, thereby identifying the uncertainty level for which the nominal regime is no longer preserved (see Supplementary Section S2).

Therefore, we combine the derived recurrence plots with a novel quantitative metric for regime preservation/identification under probabilistic parametric uncertainty, based on persistence of the recurrence pattern of the mean output signal assessed through *blob counts* (Fig. 1d, right). Our new approach, described in Fig. 4, takes as an input a recurrence plot (Fig. 4a). The recurrence plot displays distinct yellow areas, which we call *blobs*, that emerge from the background and are associated with the largest entries of  $D$ . The geometric patterns formed by such blobs distinctly characterise each given regime and we thus use them to assess regime preservation. In particular, the blobs are fundamentally different in number between the different dynamical regimes, as can be observed in Fig. 2, while their number is conserved for sufficiently small uncertainty levels, as shown in Fig. 3. Hence, the number of blobs is a reliable metric to assess the persistence of a regime, and distinguish between different regimes, in the presence of parametric uncertainty. To count the number of blobs in a recurrence plot with matrix  $D \in [0, 1]^{T \times T}$ , we introduce a threshold  $\vartheta \in (0, 1)$  and we build the Boolean version  $\mathcal{D}$  of  $D$ , such that  $\mathcal{D}_{\ell h} = 1$  if  $D_{\ell h} \geq \vartheta$  and  $\mathcal{D}_{\ell h} = 0$  otherwise. The resulting  $\mathcal{D}$  can be visualised as an image with black pixels associated with 0 entries and white pixels associated with 1 entries, as shown in Fig. 4b, c and d, where a three-dimensional visualisation of  $D$  is also provided, with the threshold represented as a grey horizontal plane.

Given the Boolean matrix  $\mathcal{D}$ , our proposed automated blob counting procedure counts the number of connected components of  $\mathcal{D}$  with value 1, corresponding to the white areas (see, e.g., Fig. 4b, c and d). As can be seen in Fig. 4b, where numerous small white dots are present, some of the connected components can be very small, and emerge due to numerical rounding errors and approximations.

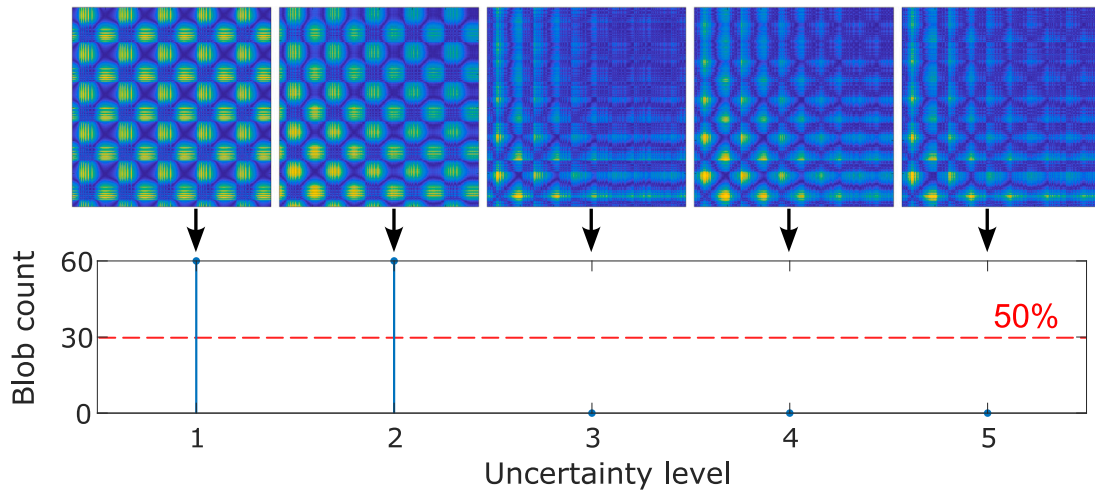
To avoid counting such small connected components, which may heavily skew the obtained blob count and its reliability, we introduce a user-defined minimum acceptable blob size  $\mathfrak{B}$  (in terms of number of pixels), so that connected components of size smaller than  $\mathfrak{B}$  are ignored in the counting process. The value of  $\mathfrak{B} = 150$  can be chosen as described in Supplementary Figure S10. Therefore, the obtained blob count, denoted henceforth by  $\mathfrak{C}$ , corresponds to the number of connected components with size at least  $\mathfrak{B}$ . As can be seen from Fig. 4b, c and d, different thresholds yield different blobs counts. Fig. 4e shows the blob count  $\mathfrak{C} = \mathfrak{C}(\vartheta)$  as a function of the threshold  $\vartheta$ . We proceed by selecting the most persistent values of  $\mathfrak{C}$ , namely, the blob counts that remain constant for the largest contiguous interval of thresholds (denoted by the green and orange horizontal lines for the two most persistent values of  $\mathfrak{C}$  in the example in Fig. 4e). We call *persistence* of a blob count the length of the largest contiguous interval of thresholds for which the blob count is exhibited. For example, in Fig. 4e, the blob count of 24 has persistence 0.58, while the blob count of 0 has persistence 0.1. This is reflected, respectively, in the green and orange bars in Fig. 4f. We thus obtain a persistence bar chart for all persistent values of  $\mathfrak{C}$ . To avoid considering blob counts with small persistence as representative of a dynamical regime, we further introduce a user-defined minimum acceptable persistence,  $\mathfrak{P}$ , which in our examples is set to be  $\mathfrak{P} = 0.05$ . Only blob counts with persistence higher than  $\mathfrak{P}$  are considered as candidate blob counts for the considered recurrence plot. For example, in Fig. 4f both blob counts of 0 and 24 have persistence larger than  $\mathfrak{P} = 0.05$ . Finally, if multiple blob count candidates have persistence larger than  $\mathfrak{P}$ , we select the blob count occurring at the lowest threshold as the actual blob count  $\mathfrak{C}^*$ . In fact, we posit that lower thresholds preserve more information of the original recurrence plot, because, as can be seen in Fig. 4d, higher thresholds lead to loss of the original geometric patterns, thereby resulting in a non-representative blob count. The same observation holds for low thresholds, which usually yield a single large connected component and, as a result, a non-informative blob count; therefore, we also discard the blob count  $\mathfrak{C} = 1$ . For example, in Fig. 4e and f, the blob count  $\mathfrak{C}^* = 24$ , which is larger than 1 and occurs at a lower threshold, is selected for the recurrence plot generated by the mean output signal of the HR model with  $b = 2.5$  and  $I \sim \mathcal{U}([3.6, 3.8])$ . If the blob counting procedure does not return any significant blob count value  $\mathfrak{C}^*$ , we do not assign any blob count to the corresponding recurrence plot, and consider that the regime has not been preserved (instead of a numerical value, we map it to a NaN; see Supplementary Figure S11 for an example).

On a side note, the most suitable threshold  $\vartheta$  to be used is not known a priori and likely depends on the specific dynamical regime. As described in Fig. 4, we drew inspiration for threshold selection from *persistent homology* [61,62], which tracks how topological structures (such as connected components or holes) form and disappear, and employs such changes to distinguish between enduring patterns, as opposed to short-lived noisy structures, in data. Similarly, we study how the blob count varies or persists through all possible thresholds  $\vartheta \in (0, 1)$ . The most persistent blob count can be different from the most frequent blob count; thus, the employed persistence bar charts differ from simple histograms.

### 2.4.3. Preservation of dynamical regimes

Consider a system of the form (1) and nominal parameters  $Z^*$ . Following the introduction of the blob counting methodology for recurrence plots in Section 2.4.2, here we present a systematic method to determine the largest support  $\mathcal{A}$  including  $Z^*$  for the uniform distribution of  $Z$  on which the nominal dynamical regime is preserved in expectation, according to the blob count metric. In particular, to identify the largest uncertainty level (i.e., the largest size of  $\mathcal{A}$ ) for which the associated recurrence plot retains a geometric pattern similar to that associated with the nominal regime, we can enlarge the size of the set  $\mathcal{A}$ , as shown in Fig. 3, as long as the resulting blob count  $\mathfrak{C}^*$  is close, in an appropriate sense, to values of the blob count that characterise persistence of the nominal regime. Hence, to assess regime preservation, we monitor persistence of the blob count as the uncertainty level is increased.

Comparing the blob count obtained for the mean output signal at each uncertainty level with the blob count obtained for the deterministic output signal with nominal parameters  $Z^*$  may be an unsuccessful approach: deterministic realisations and mean time-series can have completely different blob counts, even for small parametric uncertainties (see Supplementary Figure S12), even though their recurrence characteristics are visually similar. This discrepancy is not surprising given the different nature of these signals and the frameworks in which they are employed, which offer distinct interpretations and perspectives (see Section 1). Recall that the nominal parameter vector  $Z^*$  is located in the *interior* of a region in the parameter space that is associated with a precise dynamical regime, in the *deterministic* setting. Therefore, in the presence of parametric uncertainty, we can safely assume that, if the parametric uncertainty at the first uncertainty level is small enough (i.e.,  $\mathcal{A}$  in (1) is sufficiently small), then the corresponding blob count is representative of the average dynamical regime around the realisation at  $Z^*$ . Denote by  $\mathfrak{C}_1^*$  the blob count returned by the blob counting method in Section 2.4.2 for the smallest uncertainty level. Similarly, let  $\{\mathfrak{C}_i^*\}_{i=2}^N$  be the blob counts associated with an increasing sequence of uncertainty levels, where the subscript  $i$  grows with the uncertainty level. Let  $\gamma \in (0, 1)$  be a parameter chosen by the user. To detect loss of regime preservation, we look for the smallest index  $i > 1$  for which either  $\mathfrak{C}_i^* < \gamma \mathfrak{C}_1^*$  or  $\mathfrak{C}_i^* > (1 + \gamma)\mathfrak{C}_1^*$ . Each of these inequalities is interpreted as a significant relative change in  $\mathfrak{C}_i^*$  with respect to  $\mathfrak{C}_1^*$ . In general,  $\gamma$  may be problem-dependent and is thus added as a user-defined parameter for the method. In our examples, we use  $\gamma = 0.5$  since, in most cases, the blob count drops to zero when the uncertainty level yields a significant degradation of the recurrence pattern (as in Fig. 5); therefore,  $\gamma = 0.5$  is sufficient to exclude significant drops in blob count, while still allowing for small variations in the blob count as the uncertainty level increases. Supplementary Figure S13 investigates the suitability of this criterion in the HR model. Fig. 5 shows the sequence of blob counts, for increasing uncertainty levels, for the HR model (2) (fixing the left-end of the uncertainty intervals, as in Fig. 3). The described procedure yields the maximum tolerable uncertainty level  $i^* \leq N$  (i.e., the maximum  $i \geq 1$  such that  $\gamma \mathfrak{C}_1^* \leq \mathfrak{C}_i^* \leq (1 + \gamma)\mathfrak{C}_1^*$ ), corresponding to the maximum parametric uncertainty that the system can withstand while preserving the nominal regime, along with the maximum blob count  $\mathfrak{C}_{\max}^* = \max_{i=1, \dots, i^*} \mathfrak{C}_i^*$  associated with the preserved dynamical regime in the considered region of the parameter space.



**Fig. 5.** Blob counts with increasing uncertainty level for the recurrence plots of  $\mathbb{E}[x_1]$  for the HR model (2), with  $I = 2.8$  and  $b \sim \mathcal{U}([2.7, 2.7 + 0.03i])$ ,  $i = 1, \dots, 5$ . A significant blob count change occurs for  $i = 3$ , when the blob count drops below 50% of the first blob count (dashed red line): the maximum tolerable uncertainty level is  $i^* = 2$ .

### 2.5. Construction of the probabilistic regime preservation plot

So far, we have proposed a method, based on blob counting in recurrence plots, that allows us to determine the persistence of a regime for system (1) with parametric uncertainty affecting the nominal parameters  $Z^*$ . Consider now the case where multiple nominal parameter values  $\{Z_{(k)}^*\}_{k=1}^L$  are of interest in the parameter space  $\mathcal{P} \subset \mathbb{R}^d$  that is relevant for the model under consideration, and assume that  $\mathcal{P}$  is a hyper-rectangle. We can employ our proposed methodology to study regime preservation in the neighbourhood of each of the nominal parameter values  $\{Z_{(k)}^*\}_{k=1}^L$  in  $\mathcal{P}$ , thus assessing their local probabilistic robustness. The outcome of this analysis is a family of hyper-rectangles  $\mathcal{A}_{(k)} \subset \mathcal{P}$ ,  $k = 1, \dots, L$ , where each  $\mathcal{A}_{(k)}$  is the support of the maximum tolerable parametric uncertainty imposed on  $Z_{(k)}^*$  in (1), as defined in Section 2.4.3, and hence the sets  $\{\mathcal{A}_{(k)}\}_{k=1}^L$  indicate regions of probabilistic robustness under parametric uncertainty, according to our proposed methodology. The resulting PRP plots for the parameter space  $\mathcal{P}$  can be constructed, with two tailored approaches, in two different scenarios. In both cases, we can decide the number of uncertainty levels  $N$  and the maximum size  $\Delta Z_{\max}$  of the uncertainty interval for the parameters of interest, based on a trade-off between desired level of detail and sustainable computational effort.

**LE approach.** When no specific nominal parameter values of interest are given a priori, and we wish to gain information on probabilistic robustness properties of the regimes in the whole parameter space  $\mathcal{P}$  of interest, we determine the set of  $\{Z_{(k)}^*\}_{k=1}^L$  in  $\mathcal{P}$  with the aim of fully tiling  $\mathcal{P}$  with hyper-rectangles  $\{\mathcal{A}_{(k)}\}_{k=1}^L$ . In particular, we adopt a “left-end” (LE) approach that builds upon the same steps employed in the running example of the HR model (2) in Figs. 3 and 5. We set the nominal value  $Z_{(1)}^* \in \mathcal{P}$  as the corner of the hyper-rectangle  $\mathcal{P}$  with the smallest parameter values, and fix it as the “left-end” corner (with all entries being the smallest) of a sequence of hyper-rectangles that represent increasing parametric uncertainties. Then, we use the procedure described in Section 2.4.3 to find the blob count  $\mathcal{C}_{\max, (1)}^*$  and the maximum tolerable uncertainty level  $i_{(1)}^*$  in order to determine  $Z_{(2)}^*$  for the next iteration. We repeat the procedure until the whole  $\mathcal{P}$  has been tiled with hyper-rectangles. The LE approach yields a single PRP plot (see the examples in Fig. 1e and Fig. 6b1 and b2), in which the maximum blob counts  $\{\mathcal{C}_{\max, (k)}^*\}_{k=1}^L$  encoded by the colour of the hyper-rectangles  $\{\mathcal{A}_{(k)}\}_{k=1}^L$  convey information about the associated dynamical regimes, while the size of the hyper-rectangles, each characterised by the corresponding level of uncertainty among  $\{i_{(k)}^*\}_{k=1}^L$ , represents how much parametric uncertainty can be imposed in the neighbourhoods of the nominal parameters  $\{Z_{(k)}^*\}_{k=1}^L$  while guaranteeing preservation of the nominal regime in expectation: larger hyper-rectangles highlight regions in  $\mathcal{P}$  that exhibit greater probabilistic robustness to parametric uncertainty.

The number of hyper-rectangles and their sizes are a priori unknown. Still, the procedure is very easily implementable for a single uncertain parameter, namely if  $\mathcal{P} \subset \mathbb{R}$ , as it amounts to splitting a segment into smaller segments; examples are given in Supplementary Figure S15. For  $\mathcal{P} \subset \mathbb{R}^n$  with  $n \geq 2$ , however, the procedure is more complicated to iterate. For  $\mathcal{P} \subset \mathbb{R}^2$ , which is the parameter space relevant for the HR model (2) that we use to present our methodology, the PRP plot is easy to visualise; see Fig. 6b1 and b2. Details on the implementation of the LE approach to the HR model (2) can be found in Supplementary Figure S14.

**FG approach.** When we are given multiple nominal parameter values of interest (for instance, associated with the design of experiments), we consider a grid of equidistant points  $\{Z_{(k)}^*\}_{k=1}^L$  in  $\mathcal{P}$  that includes all the parameter choices of interest. If the chosen grid is “dense” enough, we may end up tiling the entire parameter space  $\mathcal{P}$  of interest with hyper-rectangles  $\{\mathcal{A}_{(k)}\}_{k=1}^L$ . We adopt a fixed-grid (FG) approach that considers a sequence of uncertainty intervals with increasing size, all having the nominal values  $\{Z_{(k)}^*\}_{k=1}^L$  at the centre: for each component  $j = 1, \dots, d$  of  $Z_{(k)}$  we have  $\{Z_{(k),j} \sim \mathcal{U}([Z_{(k),j}^* - 0.5(\Delta Z_{(k),j})_i, Z_{(k),j}^* + 0.5(\Delta Z_{(k),j})_i])\}_{i=1}^N$ , where  $(\Delta Z_{(k),j})_i$  is the size of the uncertainty interval for the  $j$ th component of the  $k$ th nominal parameter value at the  $i$ th uncertainty

level. The FG approach yields a pair of plots in the parameter space  $\mathcal{P}$ , as shown in the examples in Fig. 6c1 and c2 for the HR model: for each of the  $\{Z_{(k)}^*\}_{k=1}^L$ , one plot provides the corresponding blob count among  $\{c_{\max,(k)}^*\}_{k=1}^L$ , while the other plot provides the corresponding *preservation percentage*  $\{i_{(k)}^* \frac{100}{N}\}_{k=1}^L$ , where  $\{i_{(k)}^*\}_{k=1}^L$  are the maximum tolerable uncertainty levels. Since we always execute at least one iteration, the preservation percentage is always at least  $\frac{100}{N}$ .

In both cases, the PRP plots capture changes in robustness for a dynamical regime without requiring a priori knowledge on the number and types of regimes exhibited by the system.

### 3. Probabilistic robustness analysis results

We present the results of our methodology for PRA applied to systems (2) and (3), obtained with MATLAB code run on a Windows 11 Dell Inspiron 16 laptop, with 16 GB RAM and 1.90 GHz Intel i5-1340P core. We numerically integrate the systems using the function ode45 with time step  $10^{-2}$  for the HR model (2) and  $10^{-4}$  for the JR model (3), until final time 1200 for HR and 2.5 for JR, and discard an initial phase of 600 time units for HR and 1.75 time units for JR.

#### 3.1. Robustness of bursting at the single-neuron level

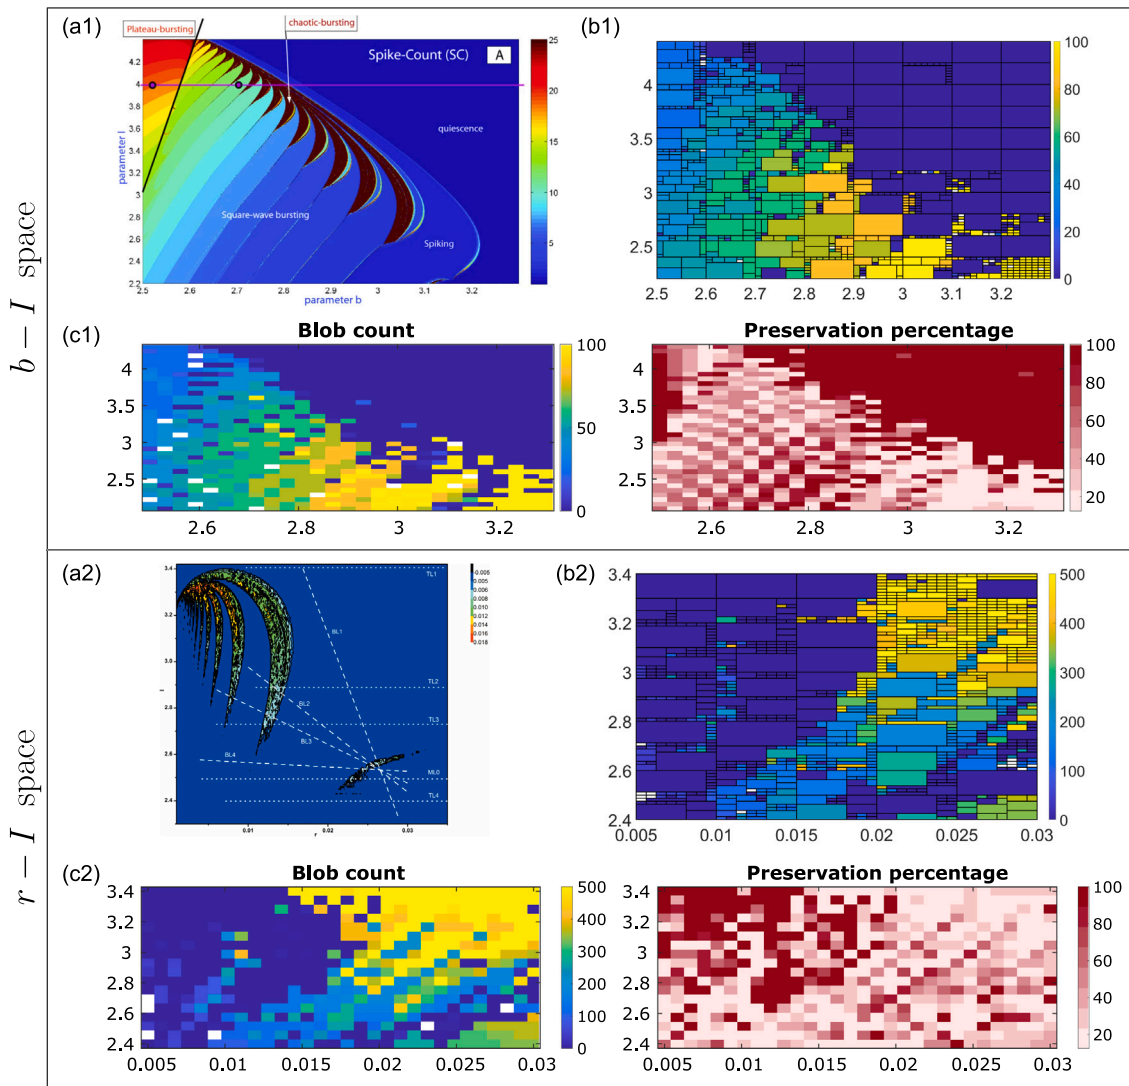
We construct PRP plots for the single-neuron HR model (2) from [37] and focus on the preservation of bursting activity in the membrane potential  $x_1(t)$  under parametric uncertainty. Our interest is driven by the fact that bursting is a fundamental element of neural computations, to reliably transmit messages to post-synaptic neurons and overcome transmission failures [63], and enable selective communication at the inter-neuron level [64]. Bursting is also important in triggering insulin secretion in  $\beta$  cells within a pancreatic islet [65] and release of growth hormone in cultured pituitary somatotrophs [66], in sensory processing in the thalamocortical relay [67], and in ensuring temporal summation and long-term potentiation [63]. So, understanding how much parametric uncertainty the system can withstand while preserving bursting is relevant for many vital biological functions. To this end, we apply our methodology, described in Section 2, to  $\mathbb{E}[x_1](t)$ . Studying mean membrane potentials is justified by the principle of averaging [14,68], a natural mechanism in which the same noisy stimulus is processed in parallel by a redundant group of neurons, whose outputs are then averaged in the post-synaptic neurons, leading to a more fault-tolerant computation. Examples of averaging include neighbouring retinal ganglion cells processing visual signals [69] and stereocilia of auditory cells processing sounds [70].

In the HR model (2), we consider three uncertain parameters:  $I$ ,  $b$  and  $r$ . Uncertainty in the input current  $I$  arises from multiple parallel neurons receiving slightly different effective inputs due to heterogeneous synaptic drive. Variability in the parameters describing neuronal dynamics is due to inter-neuron heterogeneity. We focus on the parameters  $b$ , linked to the shape of the membrane potential nullcline, and  $r$ , a time ratio that is fundamental for bursting, and consider the parameter spaces  $\mathcal{P}_1 = [b_{\min}, b_{\max}] \times [I_{\min}, I_{\max}] = [2.5, 3.3] \times [2.2, 4.4]$ , with  $r = 0.01$ , and  $\mathcal{P}_2 = [r_{\min}, r_{\max}] \times [I_{\min}, I_{\max}] = [0.005, 0.03] \times [2.4, 3.4]$ , with  $b = 3$ .

We show the PRP plots computed for  $\mathcal{P}_1$  and  $\mathcal{P}_2$  with the LE approach in Fig. 6b1 and b2, respectively, and with the FG approach in Fig. 6c1 and c2 respectively. The plots in Fig. 6a1 [21] and a2 [45] display, in the same parameter spaces as the PRP plots, the regimes resulting from different fixed parameter values in a purely *deterministic* setting, and hence offer an intrinsically different and complementary perspective. The PRP plots visualise the preservation of dynamical regimes, identified through the blob count for the recurrence plots of the mean membrane potential  $\mathbb{E}[x_1]$ , in the *probabilistic* setting of (1). The blob count is encoded either in the colour of the rectangles in Fig. 6b1 and b2, or in the colour map in the left plots in Fig. 6c1 and c2. The probabilistic robustness of nominal regimes in the presence of uniformly distributed parametric uncertainty is quantified either through the size of the rectangles in Fig. 6b1 and b2, or through the preservation percentage in the right plots in Fig. 6c1 and c2.


In the parameter space  $\mathcal{P}_1$ , the results stemming from our methodology reveal a “gradient” of colours associated with increasing blob counts, which can be observed both in Fig. 6b1 (LE approach) and Fig. 6c1, left (FG approach). Although we do not know a priori which regime corresponds to each colour, comparison with the bifurcation plot in Fig. 6a1 [21], which conveys information on the deterministic nominal regime associated with each point in the parameter space, allows us to identify the regime corresponding to a certain colour in the PRP plot: in Fig. 6b1 and Fig. 6c1, left, we can associate the upper-left light blue region with plateau bursting, the shades from turquoise to olive green to orange with square-wave bursting, and the yellow region with tonic spiking, while the dark blue area encompasses both the chaotic and the quiescence regions, as discussed in depth later in this section. The probabilistic robustness of each regime is quantified by the size of the rectangles in Fig. 6b1 (the longer the side of the rectangle, the more robust the system is with respect to the corresponding parameter) and by the preservation percentage plot in Fig. 6c1, right (the darker the colour, the higher the preservation percentage). In both cases, plateau bursting exhibits more robustness to parametric uncertainty than square-wave bursting, even though it is nominally emerging in a smaller portion of the parameter space  $\mathcal{P}_1$ . These findings are consistent with the robustness properties of bursting observed in a Hodgkin–Huxley-like model of pancreatic islets in [71].

In the parameter space  $\mathcal{P}_2$ , the deterministic plot based on Lyapunov exponents in Fig. 6a2 [45] highlights the coloured regions associated with a chaotic regime, which emerge from the blue background. In our PRP plots, these regions have zero blob count (dark blue in Fig. 6b2 and in Fig. 6c2, left) and are particularly robust to parametric uncertainty (in Fig. 6b2, many dark blue rectangles have maximal area given by the user-defined choice; in Fig. 6c2, right, the regions have high preservation percentage). However, a zero blob count also characterises the quiescence regime observed in the parameter space  $\mathcal{P}_1$ : also the quiescent dark blue areas in Fig. 6b1 and in Fig. 6c1, left, are very robust (large rectangles in Fig. 6b1 and high preservation percentage in Fig. 6c1, right). In fact, discerning meaningful patterns from mean chaotic state trajectories in the presence of parametric uncertainty is



**Fig. 6.** PRA of the bursting activity of membrane potential  $x_1$  in the HR model (2) under probabilistic parametric uncertainty, in the parameter spaces  $\mathcal{P}_1 = [b_{\min}, b_{\max}] \times [I_{\min}, I_{\max}] = [2.5, 3.3] \times [2.2, 4.4]$  with  $r = 0.01$  (top) and  $\mathcal{P}_2 = [r_{\min}, r_{\max}] \times [I_{\min}, I_{\max}] = [0.005, 0.03] \times [2.4, 3.4]$  with  $b = 3$  (bottom). Mean signals are computed through gPC surrogate models via the collocation method, with maximum expansion order  $M = 5$  and  $S_C = 250$  collocation samples. In  $\mathcal{P}_1$ ,  $\Delta b_{\max} = 0.1$  and  $\Delta I_{\max} = 0.2$ ; in  $\mathcal{P}_2$ ,  $\Delta r_{\max} = 0.005$  and  $\Delta I_{\max} = 0.1$ . In both cases,  $N = 8$  uncertainty levels are considered. In  $\mathcal{P}_1$ , we show: (a1) *deterministic* dynamical regime classification based on bifurcation analysis, using spike count as a metric (image from [21], courtesy of Springer); (b1) PRP plot obtained with the LE approach; (c1) PRP plot obtained with the FG approach, on a grid of  $27 \times 36$  equidistant points. In  $\mathcal{P}_2$ , we show: (a2) *deterministic* dynamical regime classification based on the largest Lyapunov exponent, which only highlights regions associated with the chaotic regime (image from [45], courtesy of PLoS One); (b2) PRP plot obtained with the LE approach; (c2) PRP plot obtained with the FG approach, on a grid of  $30 \times 18$  equidistant points.

particularly challenging: in a chaotic regime, the mean membrane potential computed in a probabilistic setting is indistinguishable from a zero signal affected by random noise, because it results from averaging deterministic realisations of firing patterns of the membrane potential that are very dissimilar, even for parameter values that are numerically close. This yields a rapid degradation of the recurrence patterns and hence a zero blob count, even for low uncertainty levels, thus making the mean signal in the chaotic regime indistinguishable from that in the quiescence regime, as is apparent from Fig. 6b1, where no discontinuity is visible between the chaotic bursting and the quiescence regions (for another example, see Supplementary Figure S4). On the other hand, chaos disrupts the probabilistic robustness of the bursting regimes even for small uncertainty levels, as shown by the fragmented tiling on the diagonal of Fig. 6b1, on the border between square-wave bursting and chaotic bursting regions (see also Supplementary Figure S16). Our PRP plots also reveal other, non-chaotic regions in  $\mathcal{P}_2$ , corresponding to large blob counts, where however the associated regime has a low level of robust preservation under uncertainty.



**Fig. 7.** PRA of the EEG-like signal  $x_2 - x_3$  in the JR model (3) under probabilistic parametric uncertainty, in the parameter space  $\mathcal{P}_{JR} = [A_{\min}, A_{\max}] \times [B_{\min}, B_{\max}] = [2, 12] \times [15, 31]$ , with  $N = 8$  uncertainty levels,  $\Delta A_{\max} = 2$  and  $\Delta B_{\max} = 2$ , for different values of the input  $p$  and the connectivity constant  $C$ . Mean signals are computed through gPC surrogate models via the collocation method, with maximum expansion order  $M = 5$  and  $S_C = 250$  collocation samples. PRP plots are obtained with the FG approach, on a grid of  $20 \times 20$  equidistant points.

### 3.2. Robustness at the cortical column level

The JR model (3), along with its extensions including interneurons and multiple cortical columns to better reproduce the activity of different brain regions, has been thoroughly analysed via bifurcation analysis, to investigate average inhibitory synaptic gain and time constant of mean postsynaptic potentials [18], as well as influence of input strength and dendritic time constant in a multi-input model [72], and input strength and interconnection gains in two coupled models [73]. None of these works, however, considers parametric uncertainty, which is our focus.

We study the robustness of the EEG-like signal  $y(t) = x_2(t) - x_3(t)$  with respect to parametric uncertainty in the amplitude of excitatory and inhibitory postsynaptic potentials,  $A$  and  $B$ , for the single-cortical-column JR model (3) from [38]. Fig. 7 shows the resulting PRP plots constructed, with the FG approach, in the parameter space  $\mathcal{P}_{JR} = [A_{\min}, A_{\max}] \times [B_{\min}, B_{\max}] = [2, 12] \times [15, 31]$  also considered in [38, Section 3.1], for different values of the input strength  $p$  and of the connectivity constant  $C$ . Supplementary Figure S18 provides additional PRP plots for other values of  $C$ , while PRP plots obtained with the LE approach are in Supplementary Figure S17. Moreover, Supplementary Figure S2 demonstrates the three characteristic regimes of the JR model (3): alpha-wave (highest blob count, yellow regions in Fig. 7, left), low-frequency wave (intermediate blob count, turquoise to green to orange regions in Fig. 7, left), and no-wave (zero blob count in theory, white regions in Fig. 7, left).

The PRP plots in Fig. 7 show that the alpha-wave regime is very robust to parametric uncertainty across different values of  $p$  and  $C$ , as it is mostly associated with high preservation percentages. Regions associated with the low-frequency wave regime enlarge for increasing values of  $C$ , and are on average less robust to parametric uncertainty. The no-wave regime is characterised by a constant output signal; therefore, if the considered uncertainty set is entirely contained in a no-wave region, the mean neural activity yields in

principle a zero blob count (persistent absence of recurrence pattern) that is robust to parametric uncertainty. In practice, however, some blobs may still be identified in the corresponding recurrence plot due to numerical fluctuations and due to normalisation of the distance matrix entries in (6) so that they take values in  $[0, 1]$  (see Supplementary Figure S2); however, this would result in a probabilistic robustness assessment that is not meaningful. To tackle this problem, whenever at the first uncertainty level the difference between the maximum and the minimum values of  $\mathbb{E}[y](t)$  is less than a pre-defined tolerance of  $10^{-7}$ , we assign a special value, encoded by white regions and associated with the constant no-wave regime, to both the blob count and the preservation percentage.

As mentioned earlier, zero blob counts (encoded in blue) can also arise from averaging very different deterministic realisations, which can lead to a mean signal that is indistinguishable from “random noise”; also in these cases, the condition  $\max\{\mathbb{E}[y](t)\} - \min\{\mathbb{E}[y](t)\} < 10^{-7}$  may be satisfied if the “noise” amplitude is small enough, leading to a white dot instead of a blue one. In fact, we observe that, in most PRP plots, a strip of blue dots corresponding to zero blob counts, and sometimes of white dots obtained through thresholding, clearly separates the regions associated with the alpha-wave and the low-frequency wave regimes, and also separates white regions associated with the no-wave regime from the regions associated with the other regimes; this happens because the corresponding uncertainty sets contain parameters that yield conflicting deterministic realisations, which do not allow for persistent recurrence patterns in the mean neural activity.

#### 4. Conclusions

Focusing on dynamical systems in neuroscience, we have introduced a novel methodology for PRA aimed at quantifying how much probabilistic parametric uncertainty can be imposed on the nominal system parameter values, while ensuring that the mean system output preserves the qualitative recurrence pattern associated with the nominal dynamical regime. Assessing how much uncertainty the system can withstand while still preserving the nominal regime in expectation is crucial to analyse interpretable mathematical models in neuroscience under parametric uncertainty and identify the allowable uncertainty ranges for which a life-sustaining regime is preserved. We have employed the gPC approach to efficiently calculate mean neural signals through a simpler polynomial surrogate model. The gPC method has been recently employed to generate surrogate models for uncertain systems in neuroscience [40,74,75] and systems biology [76–79], and used for global sensitivity analysis of quantities of interest in *specific* mathematical models or for parameter estimation from data. Differently from other works, here we have relied on gPC surrogate models as a key step in our proposed *general* framework for the probabilistic robustness analysis of models in neuroscience. To detect loss of regime preservation from the gPC-computed mean neural time-series, we have developed an algorithm based on the identification and the persistence of geometric patterns in the associated recurrence plots; this has been necessary to overcome the impossibility of adopting traditional neuroscientific metrics used for deterministic signals, which cannot be computed or lose their interpretability in a probabilistic setting. From a computational perspective, to systematically identify geometric patterns, we have resorted to an automated count of the number of blobs, *i.e.*, distinct regions with high values in recurrence plots; persistence of blob counts has been used as a metric for persistence of geometric patterns, and hence robust regime preservation. Finally, we have proposed two different approaches to investigate probabilistic robustness to parametric uncertainty within a whole parameter space of interest through probabilistic regime preservation (PRP) plots, which convey information both on the dynamical regimes that are preserved and on the level of tolerable parametric uncertainty affecting various nominal parameter values.

In view of its generality and flexibility, our proposed methodology is widely applicable to gain deeper insight into uncertain dynamics in neuroscience, and beyond; we have showcased its application to two fundamental models of neural dynamics.

The PRP plots for the membrane potential in the single-neuron HR model (2) reveal that plateau bursting is consistently more robust to probabilistic parametric uncertainty than square-wave bursting. Our novel methodology also provides an interesting insight: preservation of recurrence characteristics for mean state trajectories is unlikely in chaotic regions of a parameter space, as averaging completely different signals (arising even from very similar parameter values) leads to loss of any distinct geometric signature. The work [35] focuses on detection of chaos in human EEG brain signals and finds no evidence of chaotic behaviour, and our results (Fig. 6) suggest an explanation for this absence of evidence: since a single electrode measures the mean activity of several neurons, a chaotic regime would be very difficult to observe experimentally, even when some individual neurons exhibit a chaotic firing pattern as in [45].

For the EEG-like signal generated by the single-cortical-column JR model (3), we investigate the effect of uncertainty in the amplitude of excitatory and inhibitory postsynaptic potentials. PRP plots (Fig. 7) show how the regimes and the associated probabilistic robustness properties change for different values of the input and of the connectivity constant of the cortical column elements, and provide precious insight into the choice of the most appropriate parameter regions to consider when modelling the behaviour of cortical columns in a brain area of interest, based not only on the expected dynamical regimes, as in classical bifurcation analysis and in [38], but also on the observed robustness properties under parametric uncertainties, unavoidable in experimental settings.

In addition to what we have shown in this work, PRP plots can be precious to identify the most sensitive direction in which the nominal values of uncertain parameters can vary so as to transition from a dynamical regime to another, similarly to the analysis in [29]. The novel information provided by our methodology is an estimate of the robustness to parametric uncertainty of the regime to be reached, which is fundamental for the design of control actions that safely steer the neural activity towards a desired functioning. For instance, to safely design an anaesthetic action that drives brain regions to a more relaxed dynamical regime [16], we would like to minimise the changes with respect to nominal system parameter values, to avoid altering normal brain activity,

while being reasonably sure to reach an anaesthetised state that is robust to parametric uncertainty, so that the control action can be effective under perturbations.

Future research directions include refinements of the proposed PRA framework that enrich the information content of PRP plots: higher-order moments of state trajectories can be easily retrieved with gPC [80] (see, e.g., (5), where a formula for the variance of the stochastic process in (1) is also provided) and may be used to gain additional information that would, for instance, help distinguish between chaotic and quiescence regimes in a parameter space.

### CRediT authorship contribution statement

**Uros Sutulovic:** Writing – review & editing, Writing – original draft, Visualization, Software, Methodology, Investigation, Formal analysis, Data curation, Conceptualization. **Daniele Proverbio:** Writing – review & editing, Writing – original draft, Visualization, Methodology, Conceptualization. **Rami Katz:** Writing – review & editing, Writing – original draft, Visualization, Supervision, Methodology, Funding acquisition, Formal analysis, Conceptualization. **Giulia Giordano:** Writing – review & editing, Writing – original draft, Visualization, Supervision, Project administration, Methodology, Funding acquisition, Conceptualization.

### Code availability

The MATLAB code to reproduce all our results is publicly available on GitHub at <https://github.com/Uros-S/gPC-neuro-prob-robustness>.

### Funding

Work supported by the European Union through the ERC INSPIRE grant (project number 101076926). Views and opinions expressed are however those of the authors only and do not necessarily reflect those of the European Union or the European Research Council Executive Agency. Neither the European Union nor the European Research Council Executive Agency can be held responsible for them. R. Katz also acknowledges support from the Alon Fellowship, awarded by the Council of Higher Education of Israel.

### Declaration of competing interest

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests: The authors report financial support was provided by the European Union and the Council of Higher Education of Israel. The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

### Appendix A. Supplementary data

Supplementary material related to this article can be found online at <https://doi.org/10.1016/j.chaos.2026.117949>.

### Data availability

The code to reproduce all our results is publicly available on GitHub at <https://github.com/Uros-S/gPC-neuro-prob-robustness>.

### References

- [1] Kitano H. Biological robustness. *Nature Rev Genet* 2004;5(11):826–37.
- [2] Khammash M. An engineering viewpoint on biological robustness. *BMC Biol* 2016;14:22.
- [3] Blanchini F, Giordano G. Structural analysis in biology: A control-theoretic approach. *Automatica* 2021;126:109376.
- [4] Proverbio D, Katz R, Giordano G. Robustness and resilience of dynamical networks in biology and epidemiology. *Found Trends Systems Control* 2025;12(2-3):112–360.
- [5] Artime O, Grassia M, De Domenico M, Gleeson JP, Makse HA, Mangioni G, Perc M, Radicchi F. Robustness and resilience of complex networks. *Nat Rev Phys* 2024;6(2):114–31.
- [6] Gosak M, Milojević M, Duh M, Skok K, Perc M. Networks behind the morphology and structural design of living systems. *Phys Life Rev* 2022;41:1–21.
- [7] Majhi S, Ghosh S, Pal PK, Pal S, Pal TK, Ghosh D, J. Završnik M Perc. Patterns of neuronal synchrony in higher-order networks. *Phys Life Rev* 2025;52:144–70.
- [8] Bray SR, Wyss LS, Chai C, Lozada ME, Wang B. Adaptive robustness through incoherent signaling mechanisms in a regenerative brain. *Cell Rep* 2024;43(8).
- [9] Mirzakhali E, Gourgu E, Booth V, Epareanu B. Synaptic impairment and robustness of excitatory neuronal networks with different topologies. *Front Neural Circuits* 2017;11:38.
- [10] Li N, Daie K, Svoboda K, Druckmann S. Robust neuronal dynamics in premotor cortex during motor planning. *Nature* 2016;532(7600):459–64.
- [11] Aerts H, Fias W, Caeyenberghs K, Marinazzo D. Brain networks under attack: robustness properties and the impact of lesions. *Brain* 2016;139(12):3063–83.
- [12] MacLean JN, Zhang Y, Johnson BR, Harris-Warrick RM. Activity-independent homeostasis in rhythmically active neurons. *Neuron* 2003;37(1):109–20.
- [13] Koch C. *Biophysics of computation: Information processing in single neurons*. Oxford University Press; 2004.
- [14] Faisal AA, Selen LP, Wolpert DM. Noise in the nervous system. *Nature Rev Neurosci* 2008;9(4):292–303.
- [15] Hassabis D, Kumaran D, Summerfield C, Botvinick M. Neuroscience-inspired artificial intelligence. *Neuron* 2017;95(2):245–58.

- [16] Kuhlmann L, Freestone DR, Manton JH, Heyse B, Vereecke HE, Lipping T, Struys MM, Liley DT. Neural mass model-based tracking of anesthetic brain states. *NeuroImage* 2016;133:438–56.
- [17] Kerr CC, Albada SJ Van, Neymotin SA, Chadderdon GL, Robinson PA, Lytton WW. Cortical information flow in Parkinson's disease: a composite network/field model. *Front Comput Neurosci* 2013;7:39.
- [18] Subramaniam NP, Hyttinen J. Sensitivity-analysis-guided bayesian parameter estimation for neural mass models: Applications in epilepsy. *Phys Rev E* 2024;110(4):044208.
- [19] Ashourvan A, Pequito S, Khambhati AN, Mikhail F, Baldassano SN, Davis KA, Lucas TH, Vettel JM, Litt B, Pappas GJ, et al. Model-based design for seizure control by stimulation. *J Neural Eng* 2020;17(2):026009.
- [20] Izhikevich EM. *Dynamical systems in neuroscience*. MIT Press; 2007.
- [21] Barrio R, Shilnikov A. Parameter-sweeping techniques for temporal dynamics of neuronal systems: case study of Hindmarsh–Rose model. *J Math Neurosci* 2011;1:1–22.
- [22] Ashwin P, Coombes S, Nicks R. Mathematical frameworks for oscillatory network dynamics in neuroscience. *J Math Neurosci* 2016;6(1):2.
- [23] Gutenkunst RN, Waterfall JJ, Casey FP, Brown KS, Myers CR, Sethna JP. Universally sloppy parameter sensitivities in systems biology models. *PLoS Comput Biol* 2007;3(10):e189.
- [24] Goaillard J-M, Marder E. Ion channel degeneracy, variability, and covariation in neuron and circuit resilience. *Annu Rev Neurosci* 2021;44(1):335–57.
- [25] Nusser Z. Variability in the subcellular distribution of ion channels increases neuronal diversity. *Trends Neurosci* 2009;32(5):267–74.
- [26] Barmish BR. *New tools for robustness of linear systems*. Macmillan Publishing Company; 1994.
- [27] Mojtahedi M, Skupin A, Zhou J, Castaño IG, Leong-Quong RY, Chang H, Trachana K, Giuliani A, Huang S. Cell fate decision as high-dimensional critical state transition. *PLoS Biol* 2016;14(12):e2000640.
- [28] Tempo R, Calafiore G, Dabbene F. *Randomized algorithms for analysis and control of uncertain systems: with applications*, vol. 7, Springer; 2013.
- [29] Goldman MS, Golowasch J, Marder E, Abbott L. Global structure, robustness, and modulation of neuronal models. *J Neurosci* 2001;21(14):5229–38.
- [30] Foster W, Ungar L, Schwaber J. Significance of conductances in Hodgkin-Huxley models. *J Neurophysiol* 1993;70(6):2502–18.
- [31] Beer RD, Chiel HJ, Gallagher JC. Evolution and analysis of model CPGs for walking: II. general principles and individual variability. *J Comput Neurosci* 1999;7:119–47.
- [32] Marder E. Variability, compensation, and modulation in neurons and circuits. *Proc Natl Acad Sci* 2011;108(supplement\_3):15542–8.
- [33] Pouget A, Dayan P, Zemel R. Information processing with population codes. *Nature Rev Neurosci* 2000;1(2):125–32.
- [34] Golowasch J, Goldman MS, Abbott L, Marder E. Failure of averaging in the construction of a conductance-based neuron model. *J Neurophysiol* 2002;87(2):1129–31.
- [35] Perinelli A, Ricci L. Chasing chaos by improved identification of suitable embedding dimensions and lags. *Chaos: An Interdiscip J Nonlinear Sci* 2020;30(12).
- [36] Lucas M, Francois D, Mombaerts L, Donato C, Skupin A, Proverbio D. Topological analysis of brain dynamical signals indicates signatures of seizure susceptibility. *Phys Rev Res* 2025;7:043259.
- [37] Hindmarsh JL, Rose R. A model of neuronal bursting using three coupled first order differential equations. *Proc R Soc Lond [Biol]* 1984;221(1222):87–102.
- [38] Jansen BH, Rit VG. Electroencephalogram and visual evoked potential generation in a mathematical model of coupled cortical columns. *Biol Cybernet* 1995;73(4):357–66.
- [39] Xiu D, Karniadakis GE. Modeling uncertainty in flow simulations via generalized polynomial chaos. *J Comput Phys* 2003;187(1):137–67.
- [40] Sutulovic U, Proverbio D, Katz R, Giordano G. Efficient gPC-based quantification of probabilistic robustness for systems in neuroscience. In: *Proceedings of the European control conference*. 2025.
- [41] Eckmann J-P, Kamphorst SO, Ruelle D. Recurrence plots of dynamical systems. In: *Turbulence, Strange Attractors and Chaos*. World Scientific; 1995. p. 441–5.
- [42] Barmish BR, Lagoa CM. The uniform distribution: A rigorous justification for its use in robustness analysis. *Math Control Signals Systems* 1997;10(3):203–22.
- [43] Kallenberg O. *Foundations of modern probability. Probability and its Applications (New York)*. second ed. New York: Springer-Verlag; 2002.
- [44] Oksendal B. *Stochastic differential equations: An introduction with applications*. Springer Science & Business Media; 2013.
- [45] Gu H. Biological experimental observations of an unnoticed chaos as simulated by the Hindmarsh-Rose model. *PLoS One* 2013;8(12):e81759.
- [46] Cameron R, Martin W. The orthogonal development of non-linear functionals in series of Fourier–Hermite functionals. *Ann Math* 1947;385–92.
- [47] Xiu D, Karniadakis GE. The Wiener-Askey polynomial chaos for stochastic differential equations. *SIAM J Sci Comput* 2002;24(2):619–44.
- [48] Blatman G, Sudret B. Adaptive sparse polynomial chaos expansion based on least angle regression. *J Comput Phys* 2011;230(6):2345–67.
- [49] Xiu D. *Numerical methods for stochastic computations: A spectral method approach*. Princeton University Press; 2010.
- [50] Petzke F, Mesbah A, Streif S. PoCET: A polynomial chaos expansion toolbox for Matlab. In: *21st IFAC world congress*. 2020.
- [51] Jirsa VK, Stacey WC, Quilichini PP, Ivanov AI, Bernard C. On the nature of seizure dynamics. *Brain* 2014;137(8):2210–30.
- [52] Pavlov AN, Sosnovtseva OV, Mosekilde E, Anishchenko VS. Extracting dynamics from threshold-crossing interspike intervals: Possibilities and limitations. *Phys Rev E* 2000;61(5):5033.
- [53] Rosso O, Martin M, Figliola A, Keller K, Plastino A. EEG analysis using wavelet-based information tools. *J Neurosci Methods* 2006;153(2):163–82.
- [54] Meisel C, Kuehn C. Scaling effects and spatio-temporal multilevel dynamics in epileptic seizures. *PLoS One* 2012;7(2):e30371.
- [55] Lopes MA, Zhang J, Krzemiński D, Hamandi K, Chen Q, Livi L, Masuda N. Recurrence quantification analysis of dynamic brain networks. *Eur J Neurosci* 2021;53(4):1040–59.
- [56] Marwan N, Romano MC, Thiel M, Kurths J. Recurrence plots for the analysis of complex systems. *Phys Rep* 2007;438(5–6):237–329.
- [57] Prado T, Corso G, dos Santos Lima G, Budzinski R, Boaretto B, Ferrari F, Macau E, Lopes S. Maximum entropy principle in recurrence plot analysis on stochastic and chaotic systems. *Chaos: An Interdiscip J Nonlinear Sci* 2020;30(4).
- [58] Izhikevich EM. Neural excitability, spiking and bursting. *Int J Bifurc Chaos* 2000;10(06):1171–266.
- [59] Ripley BD. *Pattern recognition and neural networks*. Cambridge University Press; 2007.
- [60] Gadetsky A, Jiang Y, Brbic M. Let go of your labels with unsupervised transfer. 2024, arXiv preprint arXiv:2406.07236.
- [61] Chazal F, Michel B. *An introduction to topological data analysis: fundamental and practical aspects for data scientists*. Front Artif Intell 2021;4:667963.
- [62] Maslennikov O, Perc M, Nekorkin V. Topological features of spike trains in recurrent spiking neural networks that are trained to generate spatiotemporal patterns. *Front Comput Neurosci* 2024;18.
- [63] Lisman JE. Bursts as a unit of neural information: making unreliable synapses reliable. *Trends Neurosci* 1997;20(1):38–43.
- [64] Izhikevich EM, Desai NS, Walcott EC, Hoppensteadt FC. Bursts as a unit of neural information: selective communication via resonance. *Trends Neurosci* 2003;26(3):161–7.
- [65] Beauvois MC, Merezak C, Jonas J-C, Ravier MA, Henquin J-C, Gilon P. Glucose-induced mixed [Ca<sup>2+</sup>]<sub>i</sub> oscillations in mouse beta-cells are controlled by the membrane potential and the SERCA3 Ca<sup>2+</sup>-ATPase of the endoplasmic reticulum. *Am J Physiology-Cell Physiol* 2006;290(6):C1503–11.
- [66] Tsaneva-Atanasova K, Sherman A, van Goor F, Stojilkovic SS. Mechanism of spontaneous and receptor-controlled electrical activity in pituitary somatotrophs: experiments and theory. *J Neurophysiol* 2007;98(1):131–44.
- [67] Sherman SM, Guillery R. The role of the thalamus in the flow of information to the cortex. *Philos Trans R Soc London [Biol]* 2002;357(1428):1695–708.
- [68] van Hateren JH. A theory of maximizing sensory information. *Biol Cybernet* 1992;68(1):23–9.
- [69] Puchalla JL, Schneidman E, Harris RA, Berry MJ. Redundancy in the population code of the retina. *Neuron* 2005;46(3):493–504.

- [70] Kozlov AS, Risler T, Hudspeth A. Coherent motion of stereocilia assures the concerted gating of hair-cell transduction channels. *Nature Neurosci* 2007;10(1):87–92.
- [71] Smolen P, Rinzel J, Sherman A. Why pancreatic islets burst but single beta cells do not. Heterog Hypothesis, *Biophys J* 1993;64(6):1668–80.
- [72] Spiegler A, Kiebel SJ, Atay FM, Knösche TR. Bifurcation analysis of neural mass models: Impact of extrinsic inputs and dendritic time constants. *NeuroImage* 2010;52(3):1041–58.
- [73] Ahmadizadeh S, Karoly PJ, Nešić D, Grayden DB, Cook MJ, Soudry D, Freestone DR. Bifurcation analysis of two coupled Jansen-Rit neural mass models. *PloS One* 2018;13(3):e0192842.
- [74] Ghorri MB, Kang Y. Uncertainty quantification and sensitivity analysis of a hippocampal CA3 pyramidal neuron model under electromagnetic induction. *Nonlinear Dynam* 2023;111(14):13457–79.
- [75] Signorelli L, Manzoni A, Sætra MJ. Uncertainty quantification and sensitivity analysis of neuron models with ion concentration dynamics. *Plos One* 2024;19(5):e0303822.
- [76] Paun LM, Colebank MJ, Husmeier D. A comparison of Gaussian processes and polynomial chaos emulators in the context of haemodynamic pulse-wave propagation modelling. *Phil Trans A* 2025;383(2292):20240222.
- [77] Sumser K, Mestrom R, Tuysuz YE, Paulides MM. Exploiting polynomial chaos expansion for rapid assessment of the impact of tissue property uncertainties in low-intensity focused ultrasound stimulation. *Bioelectromagnetics* 2025;46(3):e70004.
- [78] Paulson JA, Martin-Casas M, Mesbah A. Fast uncertainty quantification for dynamic flux balance analysis using non-smooth polynomial chaos expansions. *PLoS Comput Biol* 2019;15(8):e1007308.
- [79] Hu Z, Du D, Du Y. Generalized polynomial chaos-based uncertainty quantification and propagation in multi-scale modeling of cardiac electrophysiology. *Comput Biol Med* 2018;102:57–74.
- [80] Lefebvre T. On moment estimation from polynomial chaos expansion models. *IEEE Control Syst Lett* 2020;5(5):1519–24.