

Effective stochastic models of neuroscientific data with application to weakly electric fish

Alexandre Melanson

Thesis submitted in partial fulfillment of the requirements for the Ph.D. degree in Physics

Ottawa-Carleton Institute of Physics
Department of Physics
Faculty of Science
University of Ottawa
Ottawa, Canada

Abstract

Neural systems are often stochastic, non-linear, and non-autonomous. The complex manifestation of these aspects hinders the interpretation of neuroscientific data. Neuroscience thus benefits from the inclusion of theoretical models in its methodology. Detailed biophysical models of neural systems, however, are often plagued by high-dimensional and poorly constrained parameter spaces. As an alternative, data-driven effective models can often explain the core dynamical features of a dataset with few underlying assumptions. By lumping high-dimensional fluctuations into low-dimensional stochastic terms, observed time-series can be well-represented by stochastic dynamical systems. Here, I apply this approach to two datasets from weakly electric fish.

The rate of electrosensory sampling of freely behaving fish displays spontaneous transitions between two preferred values: an active exploratory state and a resting state. I show that, over a long timescale, this rate can be modelled with a stochastic double-well system where a slow external agent modulates the relative depth of the wells. On a shorter timescale, however, fish exhibit abrupt and transient increases in sampling rate not consistent with a diffusion process. I develop and apply a novel inference method to construct a jump-diffusion process that fits the observed fluctuations. This same technique is successfully applied to intrinsic membrane voltage noise in pyramidal neurons of the primary electrosensory processing area, which display abrupt depolarization events along with diffusive fluctuations. I then characterize a novel sensory acquisition strategy whereby fish adopt a rhythmic movement pattern coupled with stochastic oscillations of their sampling rate. Lastly, in the context of differentiating between self-generated and external electrosensory signals, I model the sensory signature of communication signals between fish. This analysis provides supporting evidence for the presence of a sensory ambiguity associated with these signals.

Sommaire

Les systèmes neuronaux sont souvent stochastiques, non-linéaires, et non-autonomes. Les manifestations complexes de ces aspects compliquent l'interprétation des données neuroscientifiques. La neuroscience tire donc avantage de l'inclusion de modèles théoriques dans sa méthodologie. Les modèles biophysiques détaillés sont, toutefois, encombrés par des espaces de paramètres de haute dimension et à grandes incertitudes. De façon alternative, les modèles phénoménologiques guidés par les données peuvent souvent expliquer les aspects dynamiques principaux d'un ensemble de données en n'utilisant que quelques hypothèses de travail. En regroupant les fluctuations de haute dimension dans des termes stochastiques de basse dimension, les séries de données expérimentales peuvent être bien représentées par des systèmes dynamiques stochastiques. Ici, j'applique cette approche à deux ensembles de données portant sur les poissons faiblement électriques.

Le taux d'acquisition électrosensoriel des poissons se comportant librement démontre des transitions spontanées entre deux valeurs principales: un état d'exploration actif et un état de repos. Je montre que, sur une longue échelle de temps, ce taux peut être modélisé par un système stochastique à deux puits où un agent extérieur lent contrôle la profondeur relative des puits. Sur une échelle de temps plus courte, les poissons montrent des augmentations rapides et transitoires du taux d'échantillonnage non-conforme à un processus de diffusion. Je développe et appliquons une nouvelle méthode d'inférence pour construire un processus diffusion-saut qui reproduit les fluctuations observées. La même technique est appliquée avec succès au bruit membranaire intrinsèque des neurones pyramidaux de la zone primaire de traitement électrosensoriel, qui montrent des événements de dépolarisation soudains au sein de fluctuations diffusionnelles. Je caractérise ensuite une nouvelle stratégie d'acquisition sensorielle où les poissons adoptent des mouvements rythmiques avec des oscillations stochastiques de leur taux d'échantillonnage. Enfin, dans le contexte de la discrimination entre des signaux électrosensoriels auto-générés ou externes, je modélise la signature sensorielle des signaux de communication entre poissons. Cette analyse soutient la notion d'une ambiguïté sensorielle associée à ces signaux.

Acknowledgements

During the course of my Ph.D., I have been influenced, inspired, supported, and encouraged by many people, on many levels, without whom this thesis would not have seen the day.

As is often the case in graduate school, my research during this Ph.D. has been quite the meandering path. Yet, my supervisor, André Longtin, has been steadily patient, understanding, and generous. Importantly, he has given me the freedom to explore my own ideas, at my own pace. This independence has been instrumental in honing my creative thinking skills and in developing my own perspective on how interdisciplinary science can be done. I am also grateful for the guidance he has offered in helping me navigate the world of academia.

I have also benefited from many interactions with Prof. Leonard Maler, who has not only guided the research presented in this thesis, but has also helped me to understand the experimental aspects of neuroscience. His endless source of knowledge and his long-lasting love for science are truly inspiring to witness.

In addition, several members of the Longtin group have had a lasting influence on the way I do and think about science. In no particular order, I wish to thank Alexandre Payeur, Alexandre René, James Jun, Jorge Mejias, Wilhelm Braun, Louis Jacques, Richard Naud, Belén Sancristóbal, and Gregory Dumont. In addition, many other people have been helpful in either providing me with exciting experimental data, or with knowledge on how to interpret these data: Kevin Lee, Chadwick Boulay, Érik Harvey-Girard, Curtis Marcoux, Anh-Tuan Trinh, Nicola Jung, Avner Wallach, and Steve Clarke.

Outside academia, I have the privilege of being surrounded by a group of friends who are like family to me: Joël, Marc-André, Jean-Louis, Mathieu, Gilbert, Annik, Ariane, Marie-Christine. They have been a central part of my life and have supported me through all the good and the bad times over the past several years. I also need to thank my parents, Diane and Oscar, as they have supported me unconditionally and

have always, without hesitation, encouraged me to follow my passion, despite the highly uncertain future that this can lead to.

Lastly, the encouragement of my fiancée, Jade, were invaluable and have kept me on track even through the most difficult moments. She has taught me to take pride in what I do, and to celebrate the small victories. Her patience and cheerfulness, along with her true understanding of my motivations and goals have made the last few years much more enjoyable. On the technical side, her guru-level LaTeX skills were much appreciated when formatting this thesis.

I am immensely grateful for the funding that was granted to me during this Ph.D. from the Natural Sciences and Engineering Research Council of Canada, the Ontario Graduate Scholarship program, the University of Ottawa, and the University of Moncton.

Table of contents

Abstract	ii
Sommaire	iii
Acknowledgements	iv
Table of contents	vi
List of Figures	x
List of Tables	xii
List of Acronyms	xiii
Introduction	1
1 Background	6
1.1 Neural systems	6
1.1.1 Ion channels	6
1.1.2 Action potential	7
1.1.3 Neural networks	8
1.2 Stochastic processes	9
1.2.1 Analysis of stochastic processes	11
1.2.2 Markov processes	13
1.2.3 Diffusion processes: Langevin and Fokker-Planck equation	14
1.2.4 Jump processes: Poisson process and jump-diffusion	17
1.2.5 Numerical methods	20
1.3 Applications of effective stochastic models	23
1.4 Electric fish	26
2 Non-stationary stochastic dynamics underlie spontaneous transitions between active and inactive behavioural states	30
2.1 Introduction	31

2.2	Materials and Methods	33
2.2.1	Intermittency data	33
2.2.2	Electric Fish	34
2.2.3	Jun et al. (2014) Experiments	36
2.2.4	Data Analysis	37
	Transition Detection	41
	Active State Features	43
	State Segregation and Group Definition	44
2.2.5	Derivation of the Nonlinear Stochastic Model	44
	Integration of the Stochastic Differential Equation	49
	Estimation of the Latent Variable	51
2.3	Results	52
2.3.1	Onset-triggered Analysis	52
2.3.2	Offset-triggered Analysis	55
	Transition-Triggered PDF	60
	Numerical Integration of the Fokker-Planck Equation	61
2.3.3	Residence Time Distributions	62
	Stretched Exponential Fitting	66
2.4	Discussion	67
2.4.1	Putative Neural Structures Involved in the Transitions	68
2.4.2	Limitations of the Model	69
2.4.3	Spontaneous Movement as Decisions Making	70
2.4.4	Long tail RTDs in Intermittent Behaviour	71
2.4.5	Future Work	72
3	Data-driven jump-diffusion modelling with application to membrane voltage fluctuations in pyramidal neurons	73
3.1	Introduction	74
3.2	Methods	76
3.2.1	Definitions and overview	76
3.2.2	Choice of threshold	78
3.2.3	Jump detection	80
3.2.4	FP and true jump statistics	81
	<i>FP detection probability</i>	81
	<i>FP amplitude distribution</i>	83

	<i>True jump rate</i>	88
	<i>True jump amplitude distribution</i>	88
3.2.5	Iterative procedure, noise intensity, and drift function	89
	<i>Noise intensity</i>	90
	<i>Drift function</i>	93
3.3	Results	94
3.3.1	Validation of the FP statistics calculations	94
3.3.2	Validation of the iterative scheme	96
3.3.3	Application to experimental data	99
3.3.4	Data processing	102
3.4	Discussion	105
3.4.1	Membrane noise	106
3.5	Appendix A	108
4	Transitions between stochastic and oscillatory active sensing in pulse-type electric fish	111
4.1	Introduction	111
4.2	Results	113
4.2.1	Stochastic active sensing	115
4.2.2	Oscillatory active sensing	117
4.3	Discussion	121
4.3.1	Future work	124
4.4	Methods	125
4.4.1	Data preparation	125
4.4.2	Jump-diffusion modelling	126
4.4.3	O-state detection	127
5	Modelling the electrosensory signature of endogenous vs. exogenous chirps in electric fish	129
5.1	Introduction	129
5.2	EOD model	131
5.3	Chirp model	131
5.4	Results	134
5.5	Future work	135
5.6	Appendix A: COMSOL implementation	137

List of Figures

1	Wiener and Ornstein-Uhlenbeck process.	17
2	Jump processes.	19
3	Long-term EODR and movement data.	35
4	Active state features.	39
5	5-D scatter plot.	40
6	Heterogeneity of active states.	42
7	Non-stationary double well system.	45
8	Double well simulations - Fish A.	48
9	Onset-triggered analysis.	50
10	Double well simulation - Fish B.	51
11	Correlations between active state features.	53
12	Offset-triggered analysis.	57
13	Monte Carlo simulations.	58
14	Solution of the Fokker-Planck equation.	59
15	Offset-triggered Monte Carlo simulations.	60
16	Residence time distributions.	65
17	Jump detection.	79
18	Threshold identification.	80
19	Increment distributions.	84
20	Flowchart.	90
21	Estimation of the noise intensity.	93
22	Pure diffusion test.	96
23	Validation tests.	98
24	Estimation of the parameters of function of the SDE.	100
25	Application to experimental data.	102
26	Comparison between cells.	103
27	Dynamical regime of electrosensory sampling.	114
28	Jump-diffusion model of IPI fluctuations.	116
29	Comparison between movement and rest effective models.	117
30	Oscillatory epochs of electrosensory sampling.	119

31	O-state detection scheme.	120
32	Two types of O-states.	121
33	Autocorrelation function during O-states.	122
34	Phase space density.	123
35	Threshold for O-state detection.	128
36	Model of interacting EODs.	132
37	Chirp model	133
38	Phase reset curve.	134
39	Evidence for sensory ambiguity.	136

List of Tables

1	Numerical integration parameters for the models of both fish	51
2	Statistics of the fit for residence time distribution.	63
3	Parameters and functions for the validation cases	97
4	Comparison between estimated and correct parameters.	99
5	Comparison between rest and movement SDE parameters.	117

List of Acronyms

AM	Amplitude modulation
ACF	Autocorrelation function
CLS	Centrolateral segment
CMS	Centromedial segment
CWT	Continuous wavelet transform
EEG	Electroencephalograms
ELL	Electrosensory lateral line lobe
EOD	Electric organ discharge
EODR	Electric organ discharge rate
FP	False positive
IPI	Inter-pulse interval
PCA	Principal component analysis
PDF	Probability density function
PSD	Power spectral density
SDE	Stochastic differential equation

Introduction

“All models are wrong, but some are useful.”

- George Box

Motivation and approach

Complexity is ubiquitous in biology. Neural systems, in particular, are often stochastic, non-linear, non-autonomous, and coupled. These aspects manifest non-trivially through a range of spatial and temporal scales across the nervous system: from the stochastic activity of ion channels in neuronal membranes, to the mesoscopic oscillatory dynamics of neural networks, up to the brain-wide activity involved in spontaneous behaviour. The emergent complexity resulting from this multi-scale structure hinders the interpretation of neuroscientific data and, consequently, the formulation of unifying theories.

Although technically a branch of biology, neuroscience binds together scientific knowledge from several other fields, such as physics (e.g., electrical properties of neurons), chemistry (e.g., biochemical cascades), computer science (e.g., Big Data), and cognitive science (e.g., perception, attention) to name but a few. Like other highly interdisciplinary fields, neuroscience also includes mathematical modelling in its methodology. Research questions can thus be tackled with a combination of modelling and experiments, which allows different kinds of data to be evaluated within established theoretical frameworks. The synergy between these complimentary approaches has been instrumental in shaping our knowledge of several neural phenomena. On one hand, models can produce quantitative and falsifiable predictions, which motivates new experiments tailored to these predictions. On the other, innovative experimental techniques constantly provide new and more precise data, which calls for models with a greater level of detail and sophistication.

With the increasing amount of detail, however, mechanistic biophysical or biochem-

ical models can become plagued with high-dimensional and poorly constrained parameter spaces. If the associated parametric uncertainty is not quantified, it is then difficult to build a model that is both meaningful and scientifically valuable. As an alternative approach, complex phenomena can sometimes be understood in terms of reduced, low-dimensional representations of their core dynamics. By relying on few underlying assumptions, these so-called effective models are generally inferred directly from the data (top-down modelling), but in some cases can also be deduced from more fundamental elements (bottom-up modelling). In any case, the goal of effective modelling is limited to describing the most prominent features of a dataset with a minimal set of dynamical components. Although this descriptive nature contrasts with the predictive power of mechanistic models, sometimes a simple, self-contained description is key to gaining an intuitive understanding of a given system.

Depending on the context, effective models can have different forms. Sometimes only deterministic dynamics is needed, while in other cases, high-dimensional fluctuations are conceptualized as low-dimensional stochastic components, which are included as dynamical noise in the model. An observed time series can then be well-represented by a stochastic dynamical system, the parameters of which are inferred directly from the data. I adopt this approach here and apply it to two experimental datasets that are located at opposite ends of the spectrum of spatial and temporal scales of neural processes. The first one, at the behavioural scale, involves the temporal dynamics of sensory acquisition, while the second, at the single-neuron scale, involves intrinsic fluctuations of the membrane voltage, both of which pertain to the electrosensory system of weakly electric fish. As opposed to, e.g., electric eels that can stun their preys with electric discharges, weakly electric fish use low-voltage discharges for the purpose of navigation, communication, and prey localization. Because these fish have to expend energy to obtain information from their surroundings, this electric sense has been studied extensively in the context of active sensing.

Thesis overview and structure

I adopt here the format of a thesis by article, as opposed to the traditional format. Except for the following chapter, where I present the relevant mathematical and biological background information, all other chapters present work that is either published, submitted, or soon to be submitted to peer-reviewed journals.

In Chapter 2, I show that, over long timescales (i.e., 10 h), the rate of electrosensory sampling of freely behaving fish is bimodal and displays spontaneous transitions that can be modelled with an effective stochastic double-well system. A key component of the model is a slow external agent that modulates the relative depth of the two wells which, I hypothesize, manifests behaviourally as periods where fish favour resting over moving, and vice-versa. In addition, a conspicuous feature of the data analyzed in this chapter is the presence of abrupt and transient increases in the sensory sampling rate. These events, referred to as E-scans, were deliberately left out of this analysis, as they are not required for explaining the observed long-term behaviour. The presence of E-scans, along with the surrounding stochastic fluctuations, is reminiscent of jump-diffusion stochastic processes.

With the goal of producing an effective model for this type of dynamics, I develop in Chapter 3 a novel, generally applicable inference method that allows for a data-driven, semi-parametric estimation of the parameters and functions required to reconstruct an observed jump-diffusion process. As a test for this method, I apply it to membrane voltage fluctuations recorded in neurons of the brain region associated with electrosensory processing. These recordings display abrupt depolarization events surrounded by diffusive fluctuations, the combination of which is successfully modelled by jump-diffusion dynamics.

In Chapter 4, I apply the method developed in Chapter 3 to the same data analyzed in Chapter 2, with the goal of determining whether the short-term fluctuations in the rate of electrosensory acquisition are consistent with a jump-diffusion representation. Note, however, that this jump-diffusion inference method is only applicable to stationary data, and is thus only applied on stationary epochs of the time series analyzed in Chapter 2). When fish are resting, sensory sampling has stationary statistics and can indeed be described with an effective jump-diffusion model. During movement, however, I show that electric fish drastically change their sensory sampling strategy depending on the behavioural context. During exploratory locomotion, sensory sampling is qualitatively similar to rest and is also consistent with an effective jump-diffusion model, albeit with different parameters. At other times, however, the rate of electrosensory sampling becomes oscillatory. This occurs during what appears to be a probing behaviour, whereby fish perform rhythmic movement around a specific feature of the tank. I develop a de-

tection scheme to extract these previously undocumented oscillatory epochs from the overall time series and show that they occur with a variable frequency content.

In Chapter 5, I adapt a previously developed electrostatic model to simulate the electric fields generated by the interaction of two communicating fish. I show that voltage difference pattern sensed by the fish during communication signals is potentially ambiguous, in the sense that fish might not be able to discriminate their own signals from those emitted by other fish.

Statement of contribution

Below, I outline my contribution to the articles presented in this thesis, along with the role of all co-authors. In all cases, the research was performed under the guidance of my supervisor, Prof. André Longtin, with additional input from Prof. Leonard Maler in the Faculty of Medicine, whose laboratory provided all the data analyzed in this thesis.

- **Chapter 2:** A. Melanson, J. F. Mejias, J. J. Jun, L. Maler, & A. Longtin (2017). Non-Stationary Stochastic Dynamics Underlie Spontaneous Transitions between Active and Inactive Behavioural States. *eNeuro*, 4(2).

The data analyzed in this article were obtained and prepared by JJ, under the supervision of LM, while JM and AL provided the initial direction for the model. I further developed this initial idea and performed all subsequent data analysis and numerical simulations. I wrote the article, with the exception of a subsection of the Discussion (Putative neural structures involved in the transitions) that was written by LM. All co-authors reviewed the manuscript.

- **Chapter 3** A. Melanson and A. Longtin. Data-driven jump-diffusion modelling with application to membrane voltage fluctuations in pyramidal neurons. (under revision).

This article, initially submitted on Sept. 7th 2018, is now under revision at the *Journal of Mathematical Neuroscience* following two positive referee reports received on Jan. 7th 2019 (the submission date of this thesis). I developed the original idea and prototyped the inference method. AL suggested the application to membrane voltage data. I derived the theory, implemented it numerically, and performed all analyses and simulations. AL reviewed the manuscript.

- **Chapter 4** A. Melanson and A. Longtin. Transitions between stochastic and oscillatory active sensing in pulse-type electric fish (in preparation).

This article will be prepared for publication in the months following the submission of the thesis. I applied the inference method of Chapter 3 to the data presented in Chapter 2. I then developed the data analysis pipeline to identify and characterize oscillatory states in electrosensory sampling. I wrote the manuscript and AL reviewed it.

- **Chapter 5** A. Wallach, A. Melanson, A. Longtin, and L. Maler. Reafferent connectivity in weakly electric fish with implications for self vs non-self communication signal detection (in preparation).

AW, AL, and LM designed the research. AW performed experiments and data analysis. I implemented an existing electric organ discharge model and applied it to a simulation of two communicating fish. I performed the analysis of the results of these simulations. AW is writing the manuscript with input from AM, AL and LM

Background

In the following sections, I present a brief overview of the mathematical and biological notions that are relevant for this thesis.

1.1 Neural systems

In addition to its interdisciplinary nature, this thesis is also broad in scope, as each chapter uses a different neuroscientific context at different spatial and temporal scales. Instead of presenting a detailed review of a specific topic, I thus present here a broad, and hopefully intuitive introduction to the neural systems and phenomena that are discussed in the following chapters.

1.1.1 Ion channels

Like all cell types, neurons are enclosed by a lipid bilayer. This structure, often simply called the neuronal membrane, is composed of lipid molecules that, collectively, separate the inside of the cell from the outside environment. Importantly, the lipid bilayer is highly impermeable to ions, which are abundant in neural tissue. Most neurons contain an excess negative charge, which builds-up on the inner side of the membrane and causes

an equal build-up of positive charges on the outside. In this sense, the membrane acts as a capacitor across which one can measure a potential difference that ranges between -80 and -50 mV under normal condition (taking the outside of the cell as the reference point). This potential difference is referred to as the *resting membrane potential*, and its existence explains why the membrane is often referred to as polarized. For a cell at rest, any injection of current that reduces this potential difference (e.g., an influx of positive ions) depolarizes the neuron, while currents that increase it (e.g., the outflow of positive ions) lead to a hyperpolarization.

In order to regulate the flow of ions in and out of the cell, two special types of proteins are embedded in the membrane to allow the transport of specific ions [188]. First, *ion pumps* consume energy to move specific ions either in, or out of the cell. An example is the $\text{Na}^+\text{-K}^+$ exchanger, that pumps three Na^+ out of the cell, and two K^+ into the cell. This causes the inside of the cell to have a higher (lower) concentration of K^+ (Na^+) than the outside. In short, these pumps create and maintain ionic gradients across the cell membrane. The resulting electrochemical potential associated with each ionic species is paramount for the electrical signalling capabilities of neurons, and is exploited by the second type of membrane proteins: ion channels.

Ion channels are either open or closed. When open, an ion channel allows a specific ionic species (e.g., K^+ , Na^+ , Cl^-) to flow across the membrane according to their respective electrochemical potential, i.e., down the gradient maintained by ion pumps. These channels open and close stochastically, but the probability of finding them in a given state is a non-linear function of either the membrane potential (voltage-gated ion channels), or the local concentration of specific chemicals (ligand-gated ion channels). Even when the neuron is at rest, the collective activity of a large number of such channels opening and closing at random times creates measurable fluctuations of the membrane potential, called *channel noise*.

1.1.2 Action potential

Neurons are often referred to as input-output devices, since they receive incoming information through their dendrites (tree-like structure with several branches) and emit outgoing signals through their axon (extended cable-like structure) . The axon goes on to

make connections on the dendritic trees of other neurons at specialized structures called *synapses*. When incoming signals reach the synapse, they activate ligand-gated voltage channels that allow current to flow across the membranes. If the activation of a synapse leads to a depolarization, it is referred to as an *excitatory* synapse, whereas if it creates a hyperpolarization, it is called an *inhibitory* synapse.

The currents from all the activated synapses sum together when reaching the soma (main body of the cell). If excitatory inputs exceed inhibitory ones, a net depolarization occurs at the soma. If this depolarization exceeds a certain threshold, a stereotyped sequence of non-linear events is triggered involving voltage-gated ion channels: a large inflow of Na^+ ions rapidly depolarizes the neuron, creating a brief moment of inverted polarity (i.e., the inside is more positive than the outside). This is followed by an outflow of K^+ ions that restores the original polarity and brings the membrane potential at a slightly hyperpolarized level, after which it stabilizes back to its resting value. This event, called the *action potential*, or spike, lasts about 1 ms and propagates the whole length of the axon, where it eventually activates synapses to other neurons. Action potentials are the main communication method employed by neurons and are therefore required for information encoding and neural computations [104].

1.1.3 Neural networks

Rich dynamics can emerge when neurons are coupled together in neural networks. The behaviour of such networks depend on many parameters, e.g., their topology, the types of neurons involved, whether connections between neurons are excitatory or inhibitory, how strong the connections are (i.e., their synaptic weights), whether the network receives signals from another independent neural population (feedforward input) or signals that are themselves function of the network activity (feedback inputs), etc [36]. I describe below two examples of network dynamics that can be observed in neural systems: *bistability* and *oscillations*:

- Consider two excitatory neuronal populations, with strong internal connections, that are arranged such that they weakly excite one another, but also share a common inhibitory feedback input. The balance between cooperation and competition that arises from this structure leads to the appearance of two stable states, where

either one of the two population becomes active and the other quiets down. Interestingly, if noise is introduced in the network, transitions between these states can occur spontaneously. This situation is often referred to as winner-take-all dynamics, and is thought to be involved in decision-making, visual discrimination, and perceptual detection.

- Neuronal oscillations occur when a population of neurons displays coherent, rhythmic activity, meaning that individual neurons tend to fire spikes repetitively with some level of synchronization. Various mechanisms can lead to synchronous activity, but synaptic coupling between neurons is obviously required. To be clear, individual neurons can display rhythmic activity on their own, but collective rhythms require a network structure that allows synchronization. Such coherent activity can occur on a macroscopic scale, making them detectable through coarse measurements, such as electroencephalograms (EEG). The dynamical content of these recordings is rich, with activity appearing in several frequency bands ranging from 1 Hz (delta band), to 70 Hz and beyond (gamma band). Neuronal oscillations have been implicated in a variety of functions, such as information processing, feature binding, motor coordination attention, memory consolidation, and also in pathologies such as epilepsy.

1.2 Stochastic processes

As neural and biological systems are inherently noisy, models with probabilistic elements are common in computational neuroscience and computational biology. Below I present the mathematical basis of stochastic processes, a central topic in this thesis.

A stochastic process is a mathematical object that describes the temporal evolution of a system governed by probabilistic laws. In contrast to deterministic systems, the evolution of which is uniquely specified by its initial conditions, stochastic systems can display a variety of possible outcomes for the same initial conditions. Historical examples of stochastic processes in physics are the position of a particle that diffuses randomly in a fluid (i.e., Brownian motion), or the voltage fluctuations across a conductor caused by thermal agitation (i.e., Johnson-Nyquist noise). In fact, the study of stochastic processes was strongly influenced by the mathematical analysis of Brownian motion, initially put

forth by Einstein, Smoluchowski, and Langevin, and later formalized by Wiener, Ornstein and Uhlenbeck [62].

Formally, a stochastic process is a collection of random variables $\{X_t, t \in T\}$, indexed by a time parameter, t , that varies over a set T . As a function of time, a particular sequence of outcomes for these random variables is referred to as a *realization*, or a *sample path* of the process. That is, for each time, a specific value of X_t , say x , is drawn from a set of possible outcomes, called the state space S . Since a random variable is defined for each time, so too is a probability distribution, $P_X(x, t)$. Although this first-order distribution is useful for describing the general evolution of the process, it does not fully characterize the temporal structure imparted by the joint nature of the random variables across different times. A more complete description is provided by higher-order distributions, e.g., $P_X(x_1, \dots, x_n; t_1, \dots, t_n)$, which account for any history dependence that might govern the process. It will be shown shortly that this general description can be greatly simplified.

A stochastic process can thus be visualized as an ensemble of realizations, each of which are quantitatively different, but share a common statistical structure enforced by $P_X(x_1, \dots, x_n; t_1, \dots, t_n)$. Different types of stochastic processes exist, depending on whether S and T are discrete or continuous sets [85]. If S is discrete, for instance, X_t takes on only specific values and is referred to as a *chain*. In that case, $P_X(x_1, \dots, x_n; t_1, \dots, t_n)$ is referred to as a joint *probability mass function* and gives the probability of observing various outcomes:

$$P_X(x_1, \dots, x_n; t_1, \dots, t_n) = \text{Prob}(X_{t_1} = x_1, \dots, X_{t_n} = x_n). \quad (1)$$

On the other hand, if S is continuous, then $P_X(x_1, \dots, x_n; t_1, \dots, t_n)$ is referred to as a joint *probability density function* (PDF), defined, as usual, as the derivative of the cumulative distribution function, F_X :

$$P_X(x_1, \dots, x_n; t_1, \dots, t_n) = \frac{\partial^n F_X(x_1, \dots, x_n; t_1, \dots, t_n)}{\partial x_1 \cdots \partial x_n}, \quad (2)$$

where

$$F(x_1, \dots, x_n; t_1, \dots, t_n) = \text{Prob}(X_{t_1} \leq x_1, \dots, X_{t_n} \leq x_n). \quad (3)$$

Similarly, a process is referred to as *continuous-time*, or *discrete-time*, according to the nature of the set, T . Continuous-time stochastic processes, which are the main mathematical tool used in the following chapters, can further be classified in terms of the continuity of their trajectories: a process with continuous paths is referred to as a *diffusion process*, whereas a process with discontinuous paths, jumping around between different values and remaining constant between jumps, is called as (pure) *jump process*. A process that contains both diffusive fluctuations and the occasional jumps is, not surprisingly, referred to as a *jump-diffusion process*.

1.2.1 Analysis of stochastic processes

The statistical properties of stochastic processes can be described with the help of several functions. The mean of a process, for instance, indicates the presence of general trends:

$$\mu_X(t) = \mathbb{E}[X_t] \equiv \int_{-\infty}^{\infty} x P_X(x, t) dx, \quad (4)$$

whereas its autocorrelation gives information about temporal structures that might exist within the process:

$$R_{XX}(t, s) = \mathbb{E}[X_t X_s] \equiv \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} x_t x_s P_X(x_t, x_s; t, s) dx_t dx_s. \quad (5)$$

If the process contains any trends, it might be preferable to compare values at different times by first removing the mean, which is done in the autocovariance:

$$\begin{aligned} C_{XX}(t, s) &= \text{Cov}[X_t, X_s] \equiv \mathbb{E}\{[X_t - \mu_X(t)][X_s - \mu_X(s)]\} \\ &= R_X(t, s) - \mu_X(t)\mu_X(s), \end{aligned} \quad (6)$$

from which the familiar variance can be defined:

$$\begin{aligned} \sigma_X^2(t) &= \text{Var}[X_t] \equiv C_x(t, t) = \mathbb{E}\{[X_t - \mu_X(t)]^2\} \\ &= \mathbb{E}[X_t^2] - \mathbb{E}[X_t]^2. \end{aligned} \quad (7)$$

If the process is stationary (or, minimally, wide sense stationary), its first- and second-order distributions will be invariant to shifts of the time axis. Its mean and variance will thus be constant, $\mu_X(t) = \mu_X$ and $\sigma_X^2(t) = \sigma_X^2$, whereas its autocorrelation will be a

function of the interval, $\tau = t - s$ only:

$$R_{XX}(t, s) = R_{XX}(\tau) = E[X_t X_{t+\tau}] \quad (8)$$

In terms of terminology, it should be noted that what many authors call the *autocorrelation function* is in fact the normalized autocovariance:

$$\rho_{XX}(\tau) = \frac{E[(X_t - \mu_X)(X_{t+\tau} - \mu_X)]}{\sigma_X^2} \quad (9)$$

In the frequency domain, a stochastic process can be described by its *power spectral density*, or *power spectrum*, $S_X(\omega)$, often defined through the well-known Wiener-Khinchin relations:

$$S_{XX}(\omega) = \int_{-\infty}^{\infty} R_{XX}(\tau) e^{-i\omega\tau} d\tau \quad (10)$$

as well as

$$R_{XX}(\tau) = \frac{1}{2\pi} \int_{-\infty}^{\infty} S_{XX}(\omega) e^{i\omega\tau} d\omega. \quad (11)$$

In short, the autocorrelation and the power spectrum of a stationary stochastic process are Fourier transform pairs. It is also possible to use similar functions to compare two different stochastic processes, say X_t and Y_t . For instance, the cross-correlation is defined as

$$R_{XY}(\tau) = E[X_t Y_{t+\tau}], \quad (12)$$

and the cross-spectrum as

$$S_{XY}(\omega) = \int_{-\infty}^{\infty} R_{XY}(\tau) e^{-i\omega\tau} d\tau. \quad (13)$$

Additionally, the *spectral coherence*, or *magnitude squared coherence*, measures the linear dependence between two stochastic processes in terms of their frequency content:

$$\gamma(\omega) = \frac{|S_{XY}(\omega)|^2}{S_{XX}(\omega)S_{YY}(\omega)} \quad (14)$$

1.2.2 Markov processes

In terms of the n th-order distribution, the simplest probabilistic evolution of a stochastic process would be that of complete independence of samples at any different times: :

$$P_X(x_1, \dots, x_n; t_1, \dots, t_n) = \prod_{i=1}^n P_X(x_i, t_i). \quad (15)$$

The next most simple probabilistic evolution is that of Markov processes, which shows a wide array of possible behaviour, but are simple enough to be analyzed mathematically. The defining property of a Markov process is the so-called memoryless, or past-forgetting property, according to which the probabilistic evolution of a process is entirely determined by its current state, without any impact from its past. This is not to say that a Markov process is independent of its past, but rather, in the words of Gillespie [62], that “its future values depend on its past values only through its present values”. Mathematically, this is expressed in terms of conditional probability in the *Markov property*:

$$P_X(x_{n+1}, t_{n+1} | x_1, \dots, x_n; t_1, \dots, t_n) = P_X(x_{n+1}, t_{n+1} | x_n, t_n). \quad (16)$$

The quantity on the right-hand side of Eq. (80) is often called the *transition density*, *transition probability* (or *propagator* for an infinitesimal interval), and fully determines the evolution of a Markov process, given some initial conditions. It is therefore desirable to find methods to calculate this transition density based on the properties of the system at hand.

By applying a few fundamental rules for probabilities, the Markov property can be transformed into the cornerstone of Markov process theory, the *Chapman-Kolmogorov equation*:

$$P_X(x, t | x_0, t_0) = \int_{-\infty}^{\infty} P_X(x, t | x', t') P_X(x', t' | x_0, t_0) dx', \quad (17)$$

with $t > t' > t_0$. This imposes a condition on the transition density of all Markov processes. By transposing this condition into a more practical form, it provides a means to calculate the transition density. This can be accomplished in many ways, one of them being to consider an infinitesimal interval dt , and Taylor-expand the integrand in Eq. (69) to yield a differential equation of the form:

$$\frac{\partial}{\partial t} P_X(x, t | x_0, t_0) = \sum_{n=1}^{\infty} \frac{(-1)^n}{n!} \frac{\partial^n}{\partial x^n} [B_n(x, t) P_X(x, t | x_0, t_0)], \quad (18)$$

which is called the *Kramers-Moyal expansion*, and where the B_n 's are the *Kramers-Moyal coefficients*, often expressed in terms of the conditional moments of the increments of X_t :

$$B_n(x, t) = \lim_{dt \rightarrow 0} \frac{1}{dt} \mathbb{E} [(X_{t+dt} - X_t)^n | X_t = x]. \quad (19)$$

To develop this any further, it needs to specify whether or not the process has continuous sample paths.

1.2.3 Diffusion processes: Langevin and Fokker-Planck equation

It can be seen from Eq. (19) that knowledge of the statistics of the increments is required in order to fully characterize the transition density through Eq. (71). Fortunately, it can be shown that for a diffusion process, i.e., for a process with continuous sample paths (plus a few other mild conditions), the increments must in fact be normal random variables with mean and variance proportional to dt [62]:

$$(X_{t+dt} - X_t) \sim \mathcal{N}(A(x, t)dt, 2D(x, t)dt), \quad (20)$$

where the factor 2 is simply added for consistency with the notation of the following chapters. Because any normal random variable can be expressed in terms of a standard normal variable (i.e., mean 0 and variance 1), the increments of X_t can be written as:

$$dX_t = A(X_t, t)dt + \sqrt{2D(X_t, t)dt} \eta_t, \quad (21)$$

where η_t is distributed according to the standard normal distribution, $\eta_t \sim \mathcal{N}(0, 1), \forall t$, and $dX_t \equiv X_{t+dt} - X_t$. The functions A and D are called the *drift* and *diffusion* functions (also known as the noise intensity). Eq. (58) is an example of a *stochastic differential equation* (SDE) and provides a convenient “updating” formula that shows how the sample path of a stochastic process is influenced on one hand by a deterministic drift term, and on the other by a stochastic diffusive term. The specific form of Eq. (58) is satisfied by all continuous Markov processes [62], and serves as the mathematical backbone for the following chapters.

It is common in physics to express Eq. (58) in a form typically known as the Langevin equation:

$$\frac{dX_t}{dt} = A(X_t, t) + \sqrt{2D(X_t, t)} \Gamma_t. \quad (22)$$

The process Γ_t is called Gaussian white noise:

$$\Gamma_t \equiv \lim_{dt \rightarrow 0} \frac{\eta_t}{\sqrt{dt}} = \lim_{dt \rightarrow 0} \eta'_t, \quad (23)$$

where $\eta'_t \sim \mathcal{N}(0, 1/dt)$. The defining properties of white noise are $E[\Gamma_t] = 0$ and $E[\Gamma_t \Gamma_{t'}] = \delta(t - t')$, meaning that it has 0 mean and is uncorrelated with itself at any positive or negative lag, and has infinite variance (which yields a flat power spectrum, hence the name “white” in reference to white light containing equal power at all frequencies). Eq. (64) tells us that the formalism of SDEs in fact generalizes the notion of ordinary differential equations, which correspond to $D = 0$.

By looking back at the Kramers-Moyal expansion, Eq. (71), and by using the fact that the increments of a continuous Markov process are Gaussian, the sum can now be truncated. Indeed, the moments of the Gaussian distribution are known and can be inserted into the definition of the Kramers-Moyal coefficients. It can then be shown that only the first and the second coefficients are non-vanishing and are, respectively, $A(x, t)$ and $2D(x, t)$. The sum in Eq. (71) is thus truncated and becomes:

$$\frac{\partial P_X(x, t)}{\partial t} = -\frac{\partial}{\partial x} [A(x, t)P_X(x, t)] + \frac{\partial^2}{\partial x^2} [D(x, t)P_X(x, t)] \quad (24)$$

This partial differential equation is the well-known *Fokker-Planck equation*. With knowledge of the drift and diffusion functions, this equation can be solved for the desired transition density. Intuitively, this equation, reminiscent of the heat equation in a moving fluid, tells us how probability is moved around by the two dynamical components acting on the system, i.e., an advective (drift) and a diffusive component. It should be noted that authors often refer to $P_X(x, t|x_0, t_0)$ as “the PDF of X_T ”, and sometimes even omit the conditional.

Different functional forms for these components, i.e., $A(x, t)$ and $D(x, t)$, amount to specifying distinct stochastic processes. Here I review two of the most well-known processes:

- *Wiener process*: $A(x, t) = 0$; $D(x, t) = 1/2$

The (driftless) Wiener process (also known as Brownian motion) is the simplest continuous Markov process. It can be understood as a limiting case of the more

familiar random walk taught in standard statistical mechanics textbooks [151]. Its Fokker-Planck equation is equivalent to the one-dimensional diffusion equation, and so its transition density can be solved explicitly:

$$P(x, t|x_0, t_0) = \mathcal{N}(0, t - t_0) = \frac{1}{\sqrt{2\pi(t - t_0)}} \exp\left(-\frac{(x - x_0)^2}{2(t - t_0)}\right). \quad (25)$$

With a variance that monotonically increases with time, the Wiener process is not stationary (Fig. 1A).

In addition, its increments, dW_t , are simply standard normal random variables. For this reason, other processes are usually expressed in terms of the Wiener process, and the most common form of the SDE (58) is:

$$dX_t = A(X_t, t)dt + \sqrt{2D(X_t, t)} dW_t. \quad (26)$$

- *Ornstein-Uhlenbeck process*: $A(x, t) = -\frac{1}{\tau}x$; $D(x, t) = D$

Although it was not formalized at the time, the Ornstein-Uhlenbeck process was initially proposed by Langevin as a model for the velocity of a Brownian particle. In this sense, it describes the evolution of a particle in a parabolic potential well, i.e., under the influence of a Hookean force, in addition to being constantly knocked around by a random force. This is why it is sometimes referred to as the harmonic oscillator of stochastic analysis. Unlike the Wiener process, which is non-stationary due to its ever-increasing variance, the Ornstein-Uhlenbeck process is stable, meaning that its transition density converges to a unique, constant function of x , regardless of the initial conditions (Fig. 1B). The associated Fokker-Planck equation thus has a stationary solution, which is a Gaussian distribution in this case:

$$P_{eq}(x) = \lim_{t \rightarrow \infty} P_X(x, t|x_0, t_0) = \mathcal{N}(0, D\tau) \quad (27)$$

It should now be apparent that, with Eq. (58) and (62), the study of a stochastic process can be approached from two angles: on one hand, the SDE can be numerically integrated in order to analyze the structure of its sample paths (with, e.g., Monte Carlo methods), or, on the other, the overall evolution of the transition density can be ob-

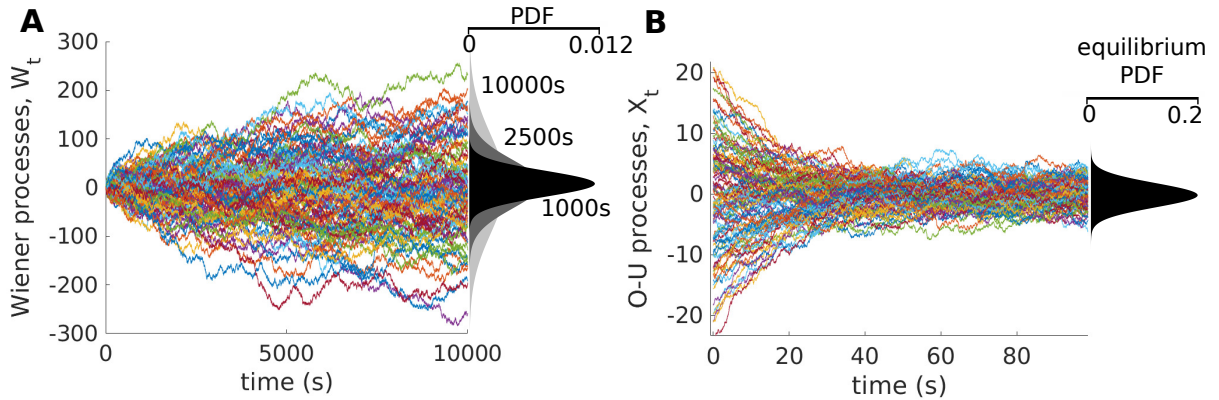


Figure 1: **A)** Several realizations of the Wiener process with $X_0 = 0$. The increasing variance is apparent from the PDFs shown at different times. **B)** Several realizations of the Ornstein-Uhlenbeck process ($\tau = 20$ s, $D = 0.2$ s $^{-1}$) with random initial conditions. The process converges to its equilibrium PDF regardless of the initial condition.

tained by solving the Fokker-Planck equation. When nonlinearities are involved, this solution will require numerical integration methods for partial differential equations (see Sec. 1.2.5).

1.2.4 Jump processes: Poisson process and jump-diffusion

Other interesting behaviour can appear if the sample paths are allowed to be discontinuous. The state space can be continuous or discrete, but in either case, the sample paths display instantaneous transitions between different values, which can occur at any random times. For this reason, the sample paths of a jump process can be entirely specified by the sequence of values visited, and the times at which the transitions occur. The times between transitions, often called *interarrival times* or *residence times*, are random variables themselves. In fact, it can be shown that, for a jump process whose parameters are constant in time, the residence times have to be exponentially distributed to ensure that the Markov property is satisfied [61].

As for the continuous case, there is a particular jump process which, due to its simplicity and generality, appears in many applications:

- *Poisson process*

The Poisson process, N_t , is a counting process, in the sense that its state space is the set of positive integers, and that its value always increases by one at every

transition, or event (Fig. 2A). With $N_0 = 0$, the sample paths of a Poisson process are staircase-like, and are fully specified by the exponentially distributed residence times. For this reason, the Poisson process is often described as a point process on the time axis. In other words, it provides time-stamps for memoryless events that occur at random times. The Poisson process can thus be used to specify the transition times of other jump processes, such as a Markov chain (Fig. 2B), or compound Poisson processes (Fig. 2C).

The increment of a Poisson process over a time interval Δt represents how many events have occurred during this time. These increments are themselves random variables and follow the Poisson distribution with parameter $\lambda\Delta t$:

$$\text{Prob}(N_{t+\Delta t} - N_t = n) = e^{-\lambda\Delta t} \frac{(\lambda\Delta t)^n}{n!}, \quad n = 0, 1, 2, \dots \quad (28)$$

where λ is called the *rate*, or *intensity*, and is the only parameter required to specify a Poisson process. One can see that, for $\lambda\Delta t \ll 1$, the probability of observing exactly one event in the interval is approximately equal to $\lambda\Delta t$.

As in the continuous case, jump processes can also be described by their transition density. The Kramers-Moyal expansion, however, is not particularly helpful here, since it can be shown that path discontinuity implies that its sum cannot be truncated like in the case of continuous processes [61]. The infinite sum in Eq. (71) is thus of limited practical use. Instead, the Chapman-Kolmogorov equation for jump processes can be expressed in integro-differential form:

$$\frac{\partial P_X(x, t|x_0, t_0)}{\partial t} = \int [W(x|x', t)P_X(x', t|x_0, t_0) - W(x'|x, t)P_X(x, t|x_0, t_0)] dx'. \quad (29)$$

Although not apparent at first glance, this equation, called the *master equation*, is more amenable to certain types of calculations than the infinite-order partial differential equation (71). The functions $W(x|x', t)$ are known as the transition rates:

$$W(x|x', t) \equiv \lim_{\Delta t \rightarrow 0} \frac{P_X(x, t + \Delta t|x', t)}{\Delta t}. \quad (30)$$

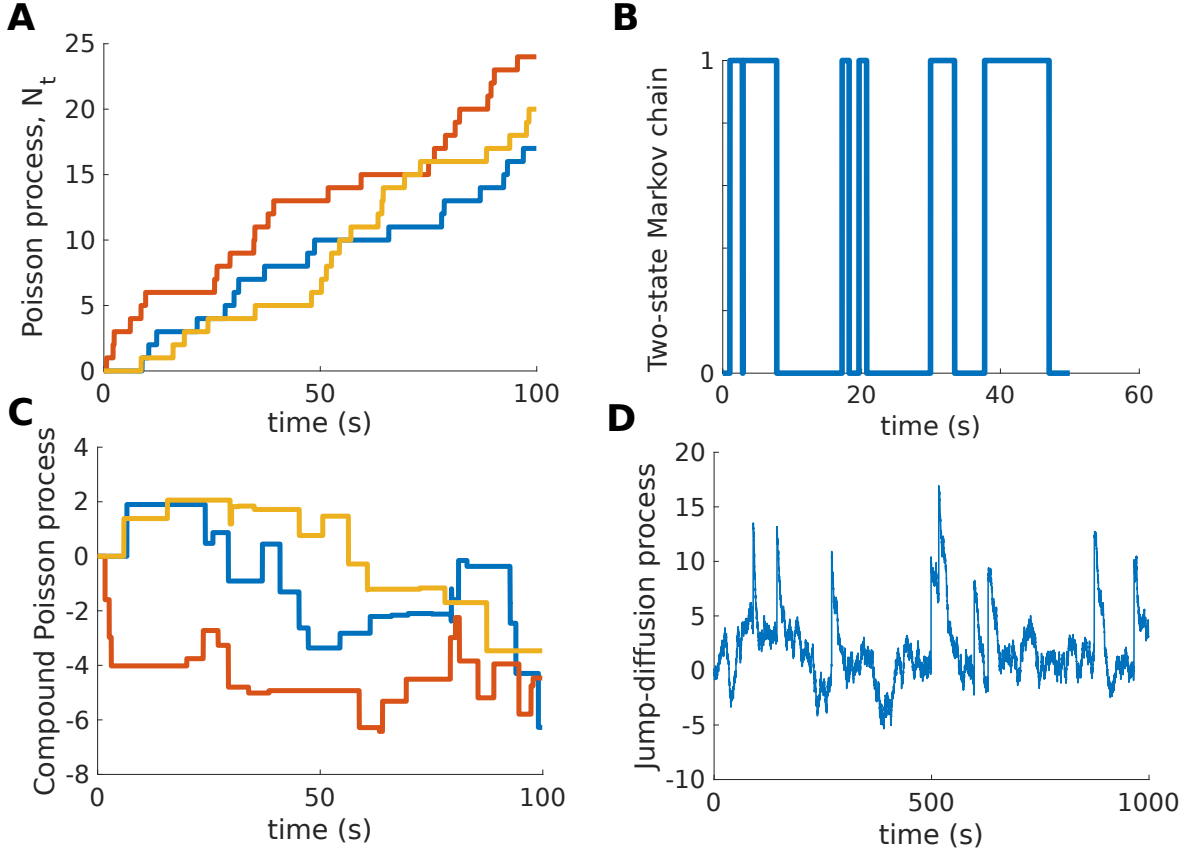


Figure 2: **A)** Realizations of the Poisson process with rate $\lambda = 0.2 \text{ s}^{-1}$. **B)** Two-state Markov chain with symmetric transition rates, whose transitions occur according to the Poisson process in A. **C)** Realizations of a compound Poisson process with the same rate as in A, but with amplitudes uniformly distributed between -2 and 2. **D)** Jump-diffusion process consisting of an Ornstein-Uhlenbeck process with $D = 0.2 \text{ s}^{-1}$ and $\tau = 20 \text{ s}$, with additional jumps of constant amplitudes $a = 10$ occurring at a Poisson rate $\lambda = 0.01 \text{ s}^{-1}$. Such processes, but with amplitudes that have their own probability distribution, will be used in Chapter 3.

Lastly, a jump process, or more precisely its increments, can be used to introduce abrupt jumps within a diffusion process. In the simplest case, the increments of a Poisson process can act as an additional stochastic term in Eq. (82) that introduces abrupt jumps of constant amplitude, a , and at a rate λ (Fig. 2D):

$$dX_t = A(X_t, t)dt + \sqrt{2D(X_t, t)} dW_t + adN_t. \quad (31)$$

Such processes will be encountered in Chapter 3. For a general jump-diffusion process, it can be shown that the time evolution of the transition density satisfies a hybrid equation

called the *differential Chapman-Kolmogorov equation* [57]:

$$\begin{aligned} \frac{\partial P_X(x, t|x_0, t_0)}{\partial t} = & -\frac{\partial}{\partial x} [A(x, t)P_X(x, t|x_0, t_0)] + \frac{\partial^2}{\partial x^2} [D(x, t)P_X(x, t|x_0, t_0)] \\ & + \int [W(x|x', t)P_X(x', t|x_0, t_0) - W(x'|x, t)P_X(x, t|x_0, t_0)] dx', \end{aligned} \quad (32)$$

which in the case of Eq. (77) reduces to

$$\begin{aligned} \frac{\partial P_X(x, t|x_0, t_0)}{\partial t} = & -\frac{\partial}{\partial x} [A(x, t)P_X(x, t|x_0, t_0)] + \frac{\partial^2}{\partial x^2} [D(x, t)P_X(x, t|x_0, t_0)] \\ & + \lambda [P_X(x - a, t|x_0, t_0) - P_X(x, t|x_0, t_0)]. \end{aligned} \quad (33)$$

1.2.5 Numerical methods

Aside from the simplest cases, such as those presented above, stochastic processes cannot be solved for analytically. Numerical methods are thus needed in order to calculate the sample paths and the transition density. Integrating SDEs such as Eq. (82) is achieved by generalizing existing methods for ordinary differential equations, such as the Euler, and Runge-Kutta methods, to accommodate the stochastic term. In contrast. the standard methods for integrating partial differential equations can be directly applied to the Fokker-Planck and the differential Chapman-Kolmogorov equations. Several methods exist, but the simplest ones are sufficient for the needs of this thesis:

- *Euler-Maruyama scheme*

For an SDE in the form of Eq. (82), let the time axis be partitioned into segments of equal duration, Δt , such that $t_{n+1} = t_n + \Delta t$, and let X_n be the value of the process at time t_n . Given an initial condition, X_0 , a sample path is obtained iteratively with:

$$X_{n+1} = X_n + A(X_n, t_n)\Delta t + \sqrt{D(X_n, t_n)}\Delta W_n \quad (34)$$

where, $\Delta W_n \sim \sqrt{\Delta t}\mathcal{N}(0, 1)$.

- *Finite volume method*

Consider the Fokker-Planck equation, Eq. (62), with $D(x, t) = D$. Let the x -axis be partitioned into a set of grid cells of equal size, Δx , such that $x_{i+1} - x_i = \Delta x$, where the x_i 's are the grid cell boundaries, or interfaces. Let $P_i(t)$ be the value of the transition density at the centre of the i^{th} cell at time t , which is delimited by

the interfaces x_{i-1} and x_i .

Let's proceed by first integrating Eq. (62) over the i^{th} grid cell:

$$\int_{x_{i-1}}^{x_i} \frac{\partial P_X(x, t|x_0, t_0)}{\partial t} dx = - \int_{x_{i-1}}^{x_i} \frac{\partial}{\partial x} [A(x, t)P_X(x, t|x_0, t_0)] dx + D \int_{x_{i-1}}^{x_i} \frac{\partial^2}{\partial x^2} P_X(x, t|x_0, t_0) dx \quad (35)$$

To handle the term on the left-hand side, a uniform profile is chosen for the transition density across the grid cell (i.e. the partial time derivative is approximated as spatially uniform): $P_X(x, t|x_0, t_0) = P_i(t) \forall x \in [x_{i-1}, x_i]$. The terms on the right-hand side are directly integrated and evaluated at the grid interfaces:

$$\frac{\partial P_i(t)}{\partial t} \Delta x = - [A(x, t)P(x, t|x_0, t_0)]_{x_i} + [A(x, t)P(x, t|x_0, t_0)]_{x_{i-1}} + D \left(\frac{\partial P(x, t|x_0, t_0)}{\partial x} \right)_{x_i} - D \left(\frac{\partial P(x, t|x_0, t_0)}{\partial x} \right)_{x_{i-1}} \quad (36)$$

The drift terms require the density to be evaluated at the grid interface. For this, I use the basic upwind scheme [150], which uses the value of P at one grid cell or another on either side of the interface depending on the sign of the drift function at that interface. This can be expressed as:

$$\begin{aligned} [A(x, t)P(x, t|x_0, t_0)]_{x_i} &= \max [A(x_i, t), 0] P_i(t) - \max [-A(x_i, t), 0] P_{i+1}(t) \\ &= \begin{cases} A(x_i, t)P_i(t) & \text{if } A(x_i, t) \geq 0 \\ A(x_i, t)P_{i+1}(t) & \text{if } A(x_i, t) < 0 \end{cases} \end{aligned} \quad (37)$$

The diffusive term, evaluated at the grid interfaces, is simply obtained with a linear profile between grid centres:

$$\left(\frac{\partial P_X(x, t|x_0, t_0)}{\partial x} \right)_{x_i} = \frac{P_{i+1}(t) - P_i(t)}{\Delta x} \quad (38)$$

For temporal discretization, I use the fully implicit scheme:

$$\int_{t_j}^{t_{j+1}} P_i(t) dt = P_i(t_{j+1}) \Delta t \quad (39)$$

which is applied by integrating Eq. (36) over one time step, Δt . Although the Crank-Nicolson scheme (which assumes a linear profile for $P_i(t)$ between t_j and t_{j+1}) is more precise for small time step, it can also produce unrealistic results if the time step is not small enough [150]. We thus err on the side of caution and instead choose the fully implicit scheme, Eq. (39). Moreover, the situations where we apply this type of numerical integration do not crucially depend on higher-order precision. Combining Eq. (39) with Eqs. (37) and (38) yields an algebraic equation involving three adjacent grid cells:

$$a_i P_{i-1}(t_{j+1}) + b_i P_i(t_{j+1}) + c_i P_{i+1}(t_{j+1}) = d_i \quad (40)$$

where,

$$\begin{aligned} a_i &= - \left\{ \max [A(x_{i-1}, t_{j+1}), 0] + \frac{D}{\Delta x} \right\} \\ b_i &= \max [A(x_i, t_{j+1}), 0] + \max [-A(x_{i-1}, t_{j+1}), 0] + \frac{2D}{\Delta x} + \frac{\Delta x}{\Delta t} \\ c_i &= - \left\{ \max [-A(x_i, t_{j+1}), 0] + \frac{D}{\Delta x} \right\} \\ d_i &= \frac{P_i(t_j) \Delta x}{\Delta t} \end{aligned}$$

Across all grid cells, this creates a system of linear equations that can be solved with the tri-diagonal matrix algorithm [150], provided that appropriate boundary conditions are enforced at both ends of the domain (i.e., absorbing or reflecting). This calculation is repeated at each time step to obtain the time-dependent transition density, $P(x, t|x_0, t_0)$. The numerical integration method outlined here works well with stationary and time-dependent parabolic partial differential equations [150], such as the Fokker-Planck equation. In Chapter 2, for instance, I apply this method to solve the Fokker-Planck equation subject to a non-stationary forcing.

1.3 Applications of effective stochastic models

Stochastic processes present a suitable framework to build effective models for the activity of many neural systems. The set of tools presented in the previous section can thus be used to examine the consequences of such models and, potentially, generate or consolidate knowledge about the original system. A challenge in this regard is to establish the correspondence between experimental time series, assumed to represent discrete samples of a stochastic process, and the theoretical quantities associated with the model. This include, notably, the estimation of the drift and diffusion function, jump rate, transition density, etc.

This approach has been successfully applied to many different situations. I review below two such examples: *spontaneous animal behaviour*, that is, the sequence of behaviours that an animal produces on its own, when there are no external perturbations, and *epileptic brain dynamics*, as measured by EEGs (i.e., electric signals produced by the brain and measured by an array of electrodes placed on the scalp). These topics were chosen with the intent of properly establishing the context and setting the stage for our study of spontaneous movement generation in electric fish (Chapter 2), and of intrinsic membrane voltage fluctuations in single neurons (Chapter 3).

- *Stochastic dynamics and behavioural states of C. elegans:*

It is commonly assumed that animal behaviour can be decomposed in a set of discrete, stereotyped templates [114], often called behavioural states. Even in simple organisms, however, identifying the neural circuits that implement these behavioural states is a major challenge. For instance, the nematode *C. elegans* (a small worm) displays a simple and well-characterized behavioural repertoire [58, 208] and its connectome¹ has been entirely mapped out [200], yet, precise neural mechanisms able to explain the range of observed behaviours remain elusive [112]. Despite this, Stephens et al. [185, 186] combine a quantitative analysis of motor output with an effective stochastic model to generate an improved understanding of *C. elegans* spontaneous behaviour. More precisely, they show that worm movements are restricted to a low-dimensional “shape space”, and they reconstruct the dynamics within this space with a two-dimensional SDE. This model

¹The “wiring diagram” for connections between all neurons of the brain.

explains several key aspects of the observed behaviour.

The instantaneous posture of a worm, at a time t , can be quantified by the tangent vector, $\hat{\boldsymbol{r}}(s, t)$, to the two-dimensional curve passing through the centre of its body, where s is the arc length along the curve. Information on the curvature as a function of s is contained in the angle corresponding to this tangent vector, $\theta(s, t)$, which is sampled at 100 equally spaced points along the length of the worm. Changes in the worm’s shape during movements can thus be quantified by a 100-dimensional, time-dependent vector. Principal component analysis of this vector shows that natural worm movements consists of a linear superposition of just four principal shapes, or “eigenworms”, $u_i(s)$, which collectively capture 95% of the shape variance: $\theta(s, t) \approx \sum_{i=1}^4 a_i(t)u_i(s)$.

Movements involving a combination of the first two modes correspond to forward and backward locomotion. In terms of the eigenworms, these movements are captured by the functions $a_1(t)$ and $a_2(t)$, which trace a ring-like structure in the (a_1, a_2) plane with a nearly constant amplitude. The dynamics in this plane is thus well-captured by a phase variable, ϕ_t . In fact, the derivative of this phase, ω_t , predicts the direction (forward or backward) and the magnitude of the worm’s velocity. The authors then reconstruct “equations of motion” in this low-dimensional (a_1, a_2) plane, which thus provide an effective description of the full, high-dimensional movement dynamics, $\theta(s, t)$ during forward or backward locomotion. Because of random fluctuations observed in the phase variable, a stochastic description of the dynamics is most natural, and here they choose a second-order system:

$$\begin{aligned} \frac{d\omega_t}{dt} &= A(\phi_t, \omega_t) + D(\phi_t, \omega_t)\Gamma_t \\ \frac{d\phi_t}{dt} &= \omega_t. \end{aligned}$$

By estimating A and D from the observations of ϕ_t and ω_t , the resulting model can be analyzed and evaluated. Notably, the deterministic phase portrait contains two limit cycles, corresponding to forward and backward motion, and two stable fixed points, corresponding both to resting, or “pause” states (there are no behavioural difference between these two stable points). The randomness imparted by the stochastic term perturbs the system and induces transitions between the

different attractors. This interplay between deterministic and stochastic dynamics helps to explain several observations, such as the residence time distribution in the moving state, the stereotypy of evoked movement in response to thermal stimuli, and changes in the movement statistics associated with foraging. Similar approaches have been applied in other systems, such as the spontaneous behaviour of fruit flies [128, 14, 123, 15].

- *Stochastic models of EEG recordings:*

EEG recordings are rich in content and change in relation to behaviour, sensory stimulation, and cognitive processes [119]. They are also valuable diagnostic tools when applied to pathological brain activity, such as epilepsy [180]. A common approach, for instance, is to calculate descriptive indices based on spectral measures or autocorrelation coefficients of EEG time series from different brain regions [9]. These indices can be compared and used to discriminate between normal and pathological neural activity, but they provide no information on the underlying dynamics. To address this, Prusseit and Lehnertz [159] reconstruct an effective Langevin equation (in the form of Eq. (19)) from EEG recordings of epileptic patients. It is thus implicitly assumed that the measured EEG signals are samples of a stochastic process X_t . They show that new indices can be obtained from empirical estimates of the drift and diffusion functions. This leads to an improved characterization of neural activity within the brain area responsible for the seizures (the epileptic focus).

Let x be the state space of X_t . In this case, x is thus the EEG signal amplitude (on the order of 10-100 μV). For stationary recordings during seizure-free intervals, the functions $A(x)$ and $D(x)$, were obtained by estimating the Kramers-Moyal coefficients, $B_1(x)$ and $B_2(x)$ (Eq. (22)), directly from the data. These coefficients were estimated for several time increments (i.e., multiples of Δt) in order to extrapolate their value as $\Delta t \rightarrow 0$. The resulting drift and diffusion functions show an overall linear and parabolic dependence on x , respectively. Interestingly, the range of $D(x)$ provides a reliable index for identifying the location of the epileptic focus.

The authors also explain that the effective Langevin model, despite being useful

for the creation of a new index, provides an incomplete representation for the dynamics of the epileptic focus. To show this, they calculate the fourth-order Kramers-Moyal coefficient, $B_4(x)$, and found it to be non-vanishing for recordings within the epileptic focus². This either indicates a deviation from the Markov property, or that discontinuities appear in the underlying process. Indeed, Anvari et al. [3] show that introducing jumps in the dynamics can account for the non-vanishing $B_4(x)$. The authors also show that for a jump-diffusion process, such as Eq. (31), with normally distributed jumps, $a \sim \mathcal{N}(0, \sigma^2)$, the following relations hold:

$$\begin{aligned} A(x) &= B_1(x) \\ D(x) + \lambda(x)\mathbb{E}[a^2] &= B_2(x) \\ \lambda(x)\mathbb{E}[a^i] &= B_i(x), \quad i > 2, \end{aligned}$$

where here, $\lambda(x)$ is a state-dependent jump rate. By estimating $B_1(x)$, $B_2(x)$, $B_4(x)$, and $B_6(x)$, a jump-diffusion model of the EEG data can be reconstructed, which provides further tools to probe the dynamics of the epileptic focus. Similar approaches have been applied to analyze heart-rate variability [153, 154], hippocampal local field potential recordings [79], the dynamics of alpha oscillations [78], and several others (reviewed in Ref. [54]).

These examples show that unobserved, high-dimensional microscopic processes, while important for the dynamical evolution and the characterization of a given system, can sometimes be modelled by low-dimensional stochastic processes, which lead to simpler descriptions than detailed mechanistic models. This can be used to circumvent a lack of knowledge about the microscopic dynamics. and improve our intuition of the overall system. This highlights the value of seeking low-dimensional representations of seemingly complex phenomena; this will be a common thread for the following chapters.

1.4 Electric fish

Understanding the human brain is obviously a main goal of neuroscience research, but studying simpler animals is often preferable and more feasible [102, 207]. As such, gen-

²According to Pawula's theorem [45], a vanishing fourth-order Kramers-Moyal coefficient, $B_4(x) = 0$, guaranties that the process is continuous, with $B_i(x) = 0$, $i > 2$.

eral principles of brain function have been extracted from studies of non-human primates [23], rodents [42], fish [99], and even insects [12]. In this thesis, I use data from weakly electric fish, a small (~ 10 cm) freshwater fish found in South America and Africa. As opposed to electric eels, rays and catfish, which can produce high voltage pulses (10-600 V) to stun their preys, weakly electric fish emit weak electric signals (< 1 V) to locate prey, navigate, and communicate. Because these fish are nocturnal and live in turbid waters [143], the electric sense is often their primary means of sensory acquisition. This sense involves two main components: generation of electric fields (*electrogenesis*), and detection of electric fields (*electroreception*).

Electric fields are produced by the dedicated *electric organ*, an elongated structure located in the tail of the fish. This organ consists of an array of cylindrical cells, called *electrocytes*, organized in a stack where they all share the same orientation. Only the end faces of each cell are excitable. An action potential elicited on the posterior (anterior) face thus generates a current of positive ions toward the head (tail) of the fish [125]. All electrocytes receive input from a hindbrain structure called the *pacemaker nucleus*, whose neurons are synchronized and fire at a regular pace [136]. This pacemaker enforces the simultaneous activation of the electrocytes, creating a net transient current across the electric organ. This yields a measurable dipolar-like electric field around the body of the fish, called the *electric organ discharge* (EOD). The exact time course, or waveform, of the EOD varies across species, but fish can be classified in two main types: *pulse-type* and *wave-type* [25].

Pulse-type fish emit brief (~ 1 ms), discrete EODs, or pulses, separated by relatively long intervals (~ 20 ms), such that the EOD rate (EODR) is relatively low, ~ 50 Hz. Although it contains a significant level of random fluctuations, the EODR is under volitional control and depends on the behavioural context [81, 97]. In Chapter 2, for instance, I analyze the bimodal nature of this EODR during intermittent locomotion³ of freely-moving fish. Although the notion of rate is often practical, an alternative, but equivalent description of electrosensory acquisition is possible in terms of the inter-pulse interval (IPI), defined as the elapsed time between two successive EODs. Even when resting, the fluctuations of the IPIs display a rich statistical structure [24]. In Chapter 4, I propose

³Alternating sequence between moving and resting states observed when unconstrained animals behave naturally [108]. Residence times in these states are often distributed according to power-laws, or other long-tail distributions [71, 158].

a stochastic dynamical representation of these fluctuations.

In contrast, wave-type fish generate a continuous, highly periodic, and quasi-sinusoidal EOD waveform at much higher rates, ~ 1000 Hz. Although wave-type fish do not control the intervals between EODs like pulse-type fish do, they still have the ability to modulate the frequency and the amplitude of the EOD oscillation, which is mostly observed in social contexts [211, 182]. For instance, fish will often produce brief (~ 20 ms) increases of their EOD frequency, called *chirps*, in response to the proximity of other fish [87]. Interaction between the EOD oscillations of two fish with different frequencies will create a beat pattern that can be altered with modulation signals such as chirps [88]. In Chapter 5, I use a computational EOD model to examine the impact of chirp signals on this beat pattern.

Regardless of their type, electric fish can detect voltage differences across their skin through an array of cutaneous *electroreceptors* distributed over their body (with a higher concentration around the head) [20]. An object in the vicinity of a fish, such as a small prey, will acquire a polarization from the EOD and will, in turn, produce a spatial pattern of voltage differences across the skin that deviates from that without the object [162, 26, 1]. This deviation, called the *electric image*, contains the relevant sensory information that is extracted by electrosensory processing areas.

The electric image is encoded and processed differently depending on the type of fish. Pulse-type fish perceive their environment in a “frame-to-frame”, or stroboscopic manner. Each pulse corresponds to a discrete sensory sample and is evaluated independently. This is evidenced by the fact that novel stimuli presented during a single EOD is sufficient to trigger the so-called novelty response, where the EODR is suddenly increased [25]. In contrast, wave-type fish do not perceive the effect of every single EOD. Instead, they extract sensory information from the slower, stimulus-induced amplitude modulation of their EOD, which then acts as a carrier wave. In either case, information from the electroreceptors are sent to pyramidal neurons⁴ of the *electrosensory lateral line lobe* (ELL), the first stage for the processing of electrosensory information. These pyramidal neurons, in turn, relay sensory information to higher brain regions, but also receive a

⁴Their name is derived from the conic, or pyramidal shape of their soma. Pyramidal neurons are ubiquitous across animals and brain regions, and are the most abundant type of excitatory neurons in the mammalian cortex [181].

significant amount of feedback, which endows the ELL with computational capabilities, such as gain control and cancellation of refferent input⁵ [20]. Given their central role in early electrosensory processing, a proper understanding of pyramidal neurons is needed to fully understand this system. In Chapter 3, I propose a jump-diffusion model to represent the intrinsic membrane voltage fluctuations observed in these neurons.

⁵Sensory input resulting from an animal's own movement.

Non-stationary stochastic dynamics underlie spontaneous transitions between active and inactive behavioural states

Abstract

The neural basis of spontaneous movement generation is a fascinating open question. Long-term monitoring of fish, swimming freely in a constant sensory environment, has revealed a sequence of behavioural states that alternate randomly and spontaneously between periods of activity and inactivity. We show that key dynamical features of this sequence are captured by a 1-D diffusion process evolving in a nonlinear double well energy landscape, in which a slow variable modulates the relative depth of the wells. This combination of stochasticity, nonlinearity, and non-stationary forcing correctly captures the vastly different timescales of fluctuations observed in the data (~ 1 to ~ 1000 seconds), and yields long-tailed residence time distributions also consistent with the data. In fact, our model provides a simple mechanism for the emergence of long-tailed distributions in spontaneous

animal behaviour. We interpret the stochastic variable of this dynamical model as a decision-like variable that, upon reaching a threshold, triggers the transition between states. Our main finding is thus the identification of a threshold crossing process as the mechanism governing spontaneous movement initiation and termination, and to infer the presence of underlying non-stationary agents. Another important outcome of our work is a dimensionality reduction scheme that allows similar segments of data to be grouped together. This is done by first extracting geometrical features in the data set and then applying principal component analysis over the feature space. Our study is novel in its ability to model non-stationary behavioural data over a wide range of timescales.

Significance Statement

Animals have the ability to initiate and terminate movement spontaneously. Given an animal moving freely in a constant sensory environment, one might expect to observe trivial behaviour. Instead, spontaneous behaviour is highly random and consists of a sequence of transitions between behavioural states. Identifying the intrinsic drivers of these transitions is necessary to understand more complex behaviours, and computational models are well suited to investigate the high-level processes governing the transitions. Here we adopt a modelling approach where the neural activity that controls movement is reduced to an effective, low-dimensional process driven by noise and evolving in a nonlinear potential landscape. We show the validity of this approach in the context of spontaneous movement initiation and termination in electric fish.

2.1 Introduction

The ability to spontaneously initiate and terminate movement is a trait shared across the animal kingdom [108, 128, 98, 11, 158]. More generally, animals can transition spontaneously and randomly through an intricate array of behavioural states, even when their sensory environment is kept constant [14]. This exemplifies how the complexity of natural animal behaviour not only stems from interactions with the environment, but also from an internally generated behavioural template. Yet, it is unknown what neural mechanisms trigger these behavioural state transitions and how randomness emerges in spontaneous behaviour. Studies where animals are relieved of sensory stimulation are thus required to isolate and understand the internal drivers of behaviour. This approach has proven useful for identifying simple principles that underly seemingly complex be-

haviour [187].

In this paper, we apply this approach by considering behavioural data, published by Jun et al. [95], from electric fish swimming freely in an empty arena. While static stimuli are always present, e.g., the tank walls, the environment is devoid of any changing sensory stimuli. Like other animals and insects in these conditions, electric fish adopt an intermittent locomotion pattern where they alternate randomly between periods of rest and periods of activity (inactive and active states, respectively). Such a binary classification of behaviour inevitably oversimplifies an animal’s larger behavioural repertoire. It does, however, underline key aspects of the intrinsic variability observed in animal behaviour. Notably, transitions between active and inactive states seem to occur spontaneously, with the time spent in each state being highly random despite the constancy of the sensory environment. These observations imply the existence of a neural control process that triggers these transitions and that imparts a high degree of stochasticity to the behavioural data. Our goal here is to use the intermittent locomotion observed in electric fish as a means to probe the dynamical origin of spontaneous movement generation. Toward this goal, we use the Jun et al. [95] data to constrain a low-dimensional, nonlinear stochastic model for the inferred neural control process.

More specifically, our goal is to identify a minimal set of dynamical components able to explain the core phenomenology of these data. We achieve this goal by developing a model where noisy fluctuations evolve as a diffusion process in a nonlinear, double well (i.e., bistable) potential landscape, and where, in addition, a latent, non-stationary deterministic forcing tilts the potential landscape back and forth on a slow timescale, thus modulating the rates of stochastic switching between the wells. This setup creates three interacting timescales of fluctuations: those within a single well (order of 1 second), those of the transitions between the wells (order of 10 to 100 seconds), and those of the latent variable (order 1000s). Note that the purpose of this model is not to explain the slowest timescale, but rather, given this slow forcing, to allow the faster fluctuations to emerge freely from it. We interpret the stochastic variable of our dynamical model as a decision-like variable that, upon crossing a threshold, triggers the transition between active and inactive states. We hypothesize that the latent variable represents slow neuromodulation that affect the animal’s internal states and its propensity to move.

Our study adds to the line of research that aims to understand complex phenomena with low-dimensional stochastic models [54]. In neuroscience, such a top-down approach has been applied to model spontaneous activity in a variety of settings [90, 33, 37, 110, 130, 78, 79]. All these studies, including the present paper, have the advantage of being data-driven and thus require very little assumptions on the underlying biophysics. Despite the simplicity of these models, they provide powerful tools for understanding the dynamical principles that govern neural processes. Our paper goes a step further in that we have developed a method to handle and model non-stationary data over long timescales.

Our main conclusion is to identify a stochastic threshold crossing process as the neural mechanism underlying the onset and offset of spontaneous movement in electric fish. In addition, we infer the existence of slow modulatory agents that impose a variable bias in the switching dynamics. The main value of our contribution is to show the applicability of a low-dimensional dynamical framework to model spontaneous natural behaviour even over long periods of time where non-stationarity is involved.

2.2 Materials and Methods

In the following, second-level subsections show the main methods and results, while third-level subsections contain the finer details and more technical information.

2.2.1 Intermittency data

Intermittent locomotion, or intermittency, has been studied extensively by biologists and ecologists [108]. These studies, however, entail monitoring unconstrained animals over long periods of time to obtain proper statistics on the movement patterns. Such experimental conditions hinder the acquisition of the concomitant neural activity, preventing any conclusions to be drawn with respect to the neural correlates of spontaneous movement. On the other hand, neuroscientists have obtained and analysed neural data pertaining to self-initiated actions, but only in the context of task-oriented experiments and only over short timescales [116, 173, 141]. These experimental paradigms, therefore, do not exactly probe natural, intrinsic animal behaviour.

This highlights the existence of a tradeoff between observing unconstrained, freely behaving animals over timescales long enough to characterize natural intermittent behaviours, and acquiring continuous neural data required to uncover the neural basis of these behaviours. The electric fish data that we use here allow us to circumvent this problem by simultaneously providing access to a movement variable and to a proxy for high-level neural activity, namely the active sensing rate of the fish, known as the electric organ discharge rate (EODR). Both of these variables can be continuously and non-invasively recorded over long timescales (4.5 - 12 hours). Over this period, the EODR creates a complex, bimodal time series that correlates strongly with the movement variable (Figure 3) and that carries information on neural activity descending from the higher brain centers [203, 60]. An adequate model for the EODR would thus give us the ability to characterize, on a phenomenological level, the neural control process underlying spontaneous movement initiation and termination in naturally behaving fish.

The modelling presented in this paper is based on data published by Jun et al. [95]. In the next section, we briefly describe the principles of electroreception in electric fish and then turn to the experimental paradigm and key results of the Jun et al. [95] experiments.

2.2.2 Electric Fish

Electric fish possess an excitable organ, the electric organ, distributed along the length of their body that can generate an electric field around them. The spatial and temporal characteristics of this electric field depend on the specific type of fish under consideration [20]. Jun et al. [95] used fish of the *gymnotus* genus, which are classified as pulse-type due to the discrete nature and to the briefness of their electric organ discharge (EOD). Each pulse creates a transient, dipolar-like electric field that can be detected by the fish through electroreceptors distributed on its skin. In the absence of surrounding objects or perturbations, each pulse creates a stereotyped spatio-temporal pattern of voltage differences on the fish's skin. This pattern is stored to memory and compared with those associated with EODs generated during normal behaviour [27]. When objects are within a fish's sensing volume, they cause deviations from this stored pattern, providing information on the type, size, and location of the object. EODs thus represent discrete sensory sampling events that allow fish to instantaneously probe their environment. This allows fish to e.g., localize and identify preys and conspecifics, as well as to navigate.

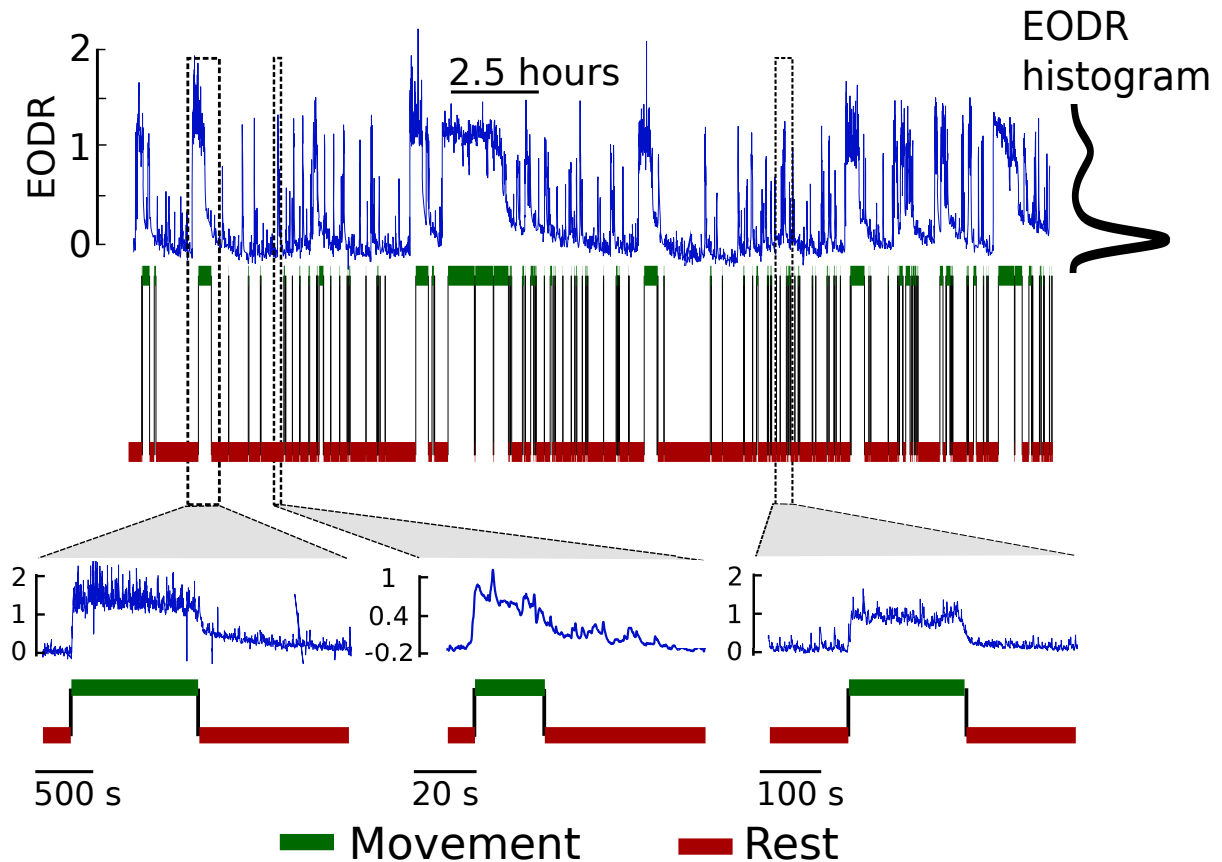


Figure 3: The EODR forms a complex, bimodally distributed time series that is highly correlated with movement. Blue traces show the EODR for Fish A, while the red and green binary traces show the movement variable, processed by a transition detection scheme (see Section 2.2.4). To allow comparison between individuals, the EODR has been rescaled as in Jun et al. [95], making it unitless. Insets show three examples of active states, and thus reveal the diversity of activity time courses in such states.

As EOD pulses are emitted in quick succession, it is often useful to consider the EODR, which provides a measure of the fish’s current level of electrosensory sampling. A high EODR corresponds to a period of heightened active sensing where the fish densely collects sensory information. Pulses are emitted at a baseline mean rate that can be modulated either spontaneously, or in response to sensory experiences. This baseline rate is established by a hindbrain pacemaker that receives input from diencephalic and medullary prepacemaker nuclei. In addition, there is evidence indicating that the EODR is strongly modulated by forebrain activity [95].

2.2.3 Jun et al. (2014) Experiments

Experiments were carried out in a featureless, circular tank surrounded by a sensory-isolation chamber that blocked external sounds, lights, and vibrations. The only stimuli available were static: the walls and floor of the tank. Animals were kept on a 12 hour light cycle and recordings were made in the dark, during the active part of the fish’s circadian rhythm. Fish were monitored, unstimulated, for long recording sessions (4.5 - 12 hours) while EODs were non-invasively recorded by an array of electrodes located on the periphery of the tank [94, 96]. The number of recording sessions varied across fish. For each session, the EODR time series were obtained by smoothing the instantaneous pulse rate followed by a rescaling to allow comparison across individuals (all EODR traces shown here are thus unitless). The rescaling mapped the median EODR values during inactive and active states to 0 and 1, respectively. The time series from each sessions were then stitched together to obtain what we refer to as the pooled EODR time series.

To avoid the need for large data storage space, the movement levels of fish were obtained directly from the recorded EODs rather than from video recordings [94, 96]. Movement information is conveyed by the EOD amplitudes at the different recording electrodes. Jun et al. [95] showed that fluctuations of these amplitudes correlate with fish movement and can thus be converted into a movement variable. This experimental paradigm was applied to five fish of unknown sex for a total of 207 hours of recording. For the analyses and modelling carried out in this paper, we showcase only two of these fish, which we refer to as Fish A and Fish B. To show the general applicability of the model that we propose below, we choose these fish because they differ the most in terms of their behavioural data and thus span the largest range of behavioural types. Fish A (7 recording sessions) is an older fish that spent most of its time in the inactive state, while Fish B (8 recording sessions) is younger and much more active, with more transitions between active and inactive states. The data from the three other fish not analysed here are qualitatively similar to those of Fish B. These fish would thus be modelled in the same way as Fish B.

These experiments revealed a strong correlation between the movement variable and the EODR: movement onset is always concomitant with a significant increases of the EODR. Moreover, the EODR tends to adopt two preferred values, as shown by the EODR histogram (Figure 3), with increased fluctuations associated with the elevated

value. By visual inspection of the EODR time series, one sees that its evolution resembles diffusional trajectories around two stable states, interrupted by sharp transitions between these states. This observation provides the foundation of the model presented below.

The fact that the EODR is bimodally distributed and that the elevated EODR value is coupled to movement allows for the definition of two behavioural attractor states, namely the active and inactive states mentioned above. In the case of these fish, active states are thus periods of movement with an elevated active sensing rate, and the transition from inactive-to-active state is concomitant with movement onset. In addition, Jun et al. [95] report a delayed correlation on a shorter timescale, with an upward transition in EODR preceding that in movement by 1 - 4 seconds. Note that this preparatory ramp-up of the EODR is not meant to be captured by the model presented here. These observations lead to the hypothesis that a single neural control process is responsible for coupling the EODR to movement and, importantly, for triggering transitions between the behavioural attractor states [95]. Because increases in EODR must be due to increased neural activity of neurons providing descending input to the pacemaker nucleus, information on the neural control process must be contained in the EODR time series. Our modelling approach and interpretations are predicated on exploiting the information content of this time series.

2.2.4 Data Analysis

We now turn to the details of the data analysis scheme that we develop to further probe the Jun et al. [95] data. This is followed by the derivation and applications of the non-linear stochastic model we propose for these data.

By visually inspecting the EODR time series, one can observe that the active states are not all alike (Figure 3, insets). Not only do they span three orders of magnitude in terms of duration (10 - 1000 seconds approximately), but they also cover a range of different shapes; this is particularly striking in the case of Fish A. For instance, there is variability in the height of the jump around inactive-to-active transitions, which we refer to as the transition amplitude. In addition, active states show, with a varying degree, a slow tail-like decay after movement offset. We implicitly assume that if active states

have similar geometrical attributes, i.e. if they "look" the same, then they are probably generated by similar neural dynamics. However, the aforementioned variability in the shape of active states leads us to hypothesize that they are not all generated by the same underlying dynamical template. To appropriately study the EODR time series, it is then necessary to group similar active states together, so as to avoid erroneously comparing states that are potentially of different neural origin, i.e., active states where the associated neural control process operates in a different dynamical regime.

To achieve this goal, we develop a dimensionality reduction scheme, similar in principle to spike sorting [115], where active states are sorted with respect to their geometrical similarities using principal components analysis (PCA, a statistical procedure by which orthogonal directions of maximal variance are identified in the data space). Although it is possible to apply PCA directly to the active state traces, we achieve better separation and grouping by instead quantifying active states with a small number of geometrical features, and then applying PCA on this feature space. The first step consists of breaking the EODR time series in a sequence of active and inactive states. To do so, the movement time series is piped through a transition detection scheme that assigns a time stamp to each transition (see Section 2.2.4). An active state is then defined as a segment of the EODR time series located between a movement onset time and a movement offset time.

The second step consists of assigning a set of geometrical features for each active state (see Figure 4). Based on preliminary exploration of the data, we choose five such features, referred to as f_1 to f_5 : the transition amplitude, the duration of the state, the average EODR during the state, the variance of the EODR during the state, and the duration of the decaying tail after movement offset. Although one might expect the transition amplitude (f_1) and the average EODR (f_3) to be equal, they in fact often differ. These specific features are chosen because they show significant variability across active states. Also recall that, although the EODR has been rescaled to show transitions between 0 and 1, this rescaling does not remove the variability observed in f_3 : to achieve this rescaling, we first obtain the median EODR for all inactive states combined, and for all active states combined, then it is these median values that are mapped to 0 and 1. The result is a rescaled EODR, the fluctuations of which yield median values in the inactive and active states of 0 and 1. The individual active states, however, still show

an average value that varies around 1.




Figure 4: Active states are characterized by five geometrical features, f_1 to f_5 , (see Section 2.2.4) that allow the dimensionality of the EODR time series (blue trace) to be reduced. t_{on} and t_{off} are the movement onset and offset times, respectively.

With each active state represented by five coordinates, the entire time series can be visualized as a cluster of points in 5-D space, each point representing a single active state (Figure 5A). This process reveals some unexpected correlations between all of the defined features. For instance, the transition amplitude (f_1) correlates positively with the duration of the state (f_2) (see Section 2.3.1) and with the duration of the decaying tail after movement offset (f_5). This suggests that the structure and the duration of the active state is, to some extent, predetermined right from the onset. This hints at an underlying factor controlling the active states, i.e. that they are not realized from purely random processes. A positive correlation is also observed between the mean (f_3) and variance (f_4) of the EODR during active states. Assuming a common cause for these correlations, we apply PCA on this 5-D cluster of points in the hope of extracting some new variables (i.e., the first few principal components) along which the active states would be

sorted according to their geometrical similarities. We find that, in fact, the first principal component (PC1) by itself fulfils this role satisfactorily. Although this type of sorting scheme usually leads to a search for clusters, as for spike sorting for instance, no such clusters are to be found in our case. Instead, active states are distributed continuously along the PC1 axis (Figure 5B). What we observe, however, is that active states that are neighbours along this axis do look similar to one another, which is the main reason why this approach is useful here.

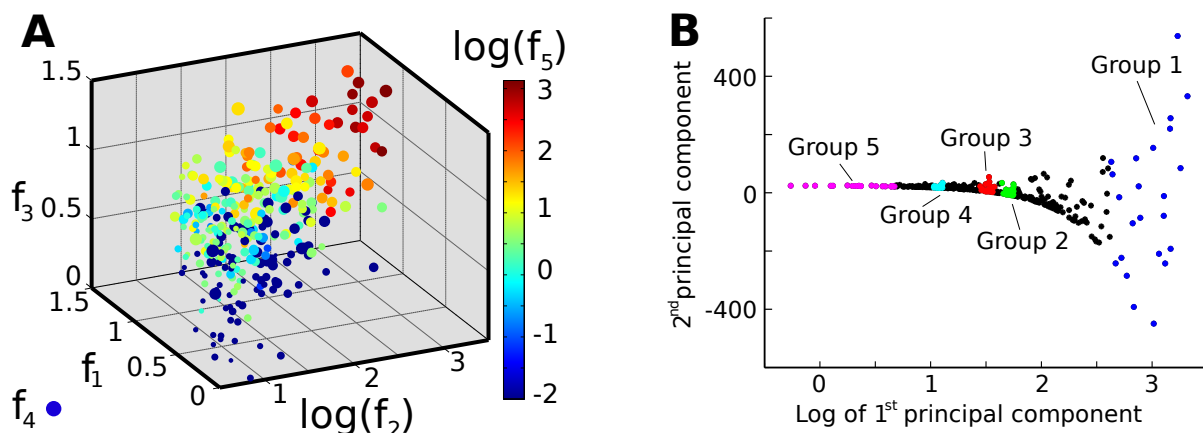


Figure 5: Correlations between features are revealed once the data set is cast onto feature space. This makes PCA useful for segregating states of similar shape. **A)** 5-D scatter plot of the active state features from Fish A. Each point corresponds to a single active state. The size of the points represent feature 4, and the colours, feature 5. **B)** Same scatter as in A, but projected onto the plane of the first two principal components. Again, each point (black and coloured) represents an active state, and colours are only added for visual representation of the different groups (colours are unrelated to those in A).

Indeed, this observation provides us with a tool to visualize and analyse the data in a new way: we can now easily plot the traces of several similar-looking active states on top of each other, and do so for the many different shapes of active states that populate the data set. We do this by grouping together active states that are neighbours along the PC1 axis, each group centered on a different location on this axis (Figure 5B). As there are no real boundaries between the active states in PC space, the number of groups we choose to define is arbitrary. Yet, we find that defining 5 distinct groups is enough to adequately sample the spectrum of different active state shapes found along the PC1 axis. We also find that having 20 active states per group is large enough to calculate representative average traces, yet few enough such as to avoid grouping differently shaped

states together (Figure 6). Note that, for Fish B, although the active states of Groups 1 to 4 look similar to one another, they do indeed differ significantly in terms of their total duration and decay duration (which is not shown in Figure 6 since the traces are aligned with respect to movement onset). This similarity is largely explained by the fact that, as opposed to Fish A, the transition amplitude does not vary consistently across groups. By inspecting the left panel of Figure 6 more closely, one might also notice that the slope of the average EODR varies across groups. This means that this slope could also have been used as a geometric feature defining each active state. However, due to the high level of fluctuations in the active state, rigorously extracting this value for individual traces is problematic and, therefore, this slope was not considered in the feature space.

In summary, each group represents a stereotypical shape of active state that is found all across the data set. For instance, for Fish A, Group 1 consists of states with high values for the PC1, which corresponds to those states with long durations, long decaying tails, and higher mean values for the EODR. States that belong to Group 5, on the other hand, have low PC1 values and have short durations with low transition amplitudes. Moreover, because all active states possess time stamps for movement onset and offset times, we have the freedom to analyse the groups by aligning their states with respect to either of these times. In Figure 6, for instance, the movement onset time was used as reference to align the active states. Note that, since active states within a group have slightly different durations, aligning them with respect to movement onset causes a de-synchronization around the movement offset time, i.e., the active states do not all drop off at the same time. The ability to extract and group similar active states together is essential for the modelling carried out in this paper. In addition, this analysis pipeline could be applied to other types of time series and constitute a useful approach to identify and handle suspected non-stationary elements in stochastic data.

Transition Detection

To assign a time stamp to movement onset and offset times, we start by compiling the histogram for the movement variable, which appears unambiguously bimodal, and then obtain the values for the local minimum as well as the two adjacent local maxima. To remove undesired rapid fluctuations in the movement variable, we smooth it with a moving average filter (window size: 1 seconds) and obtain the transition times from this

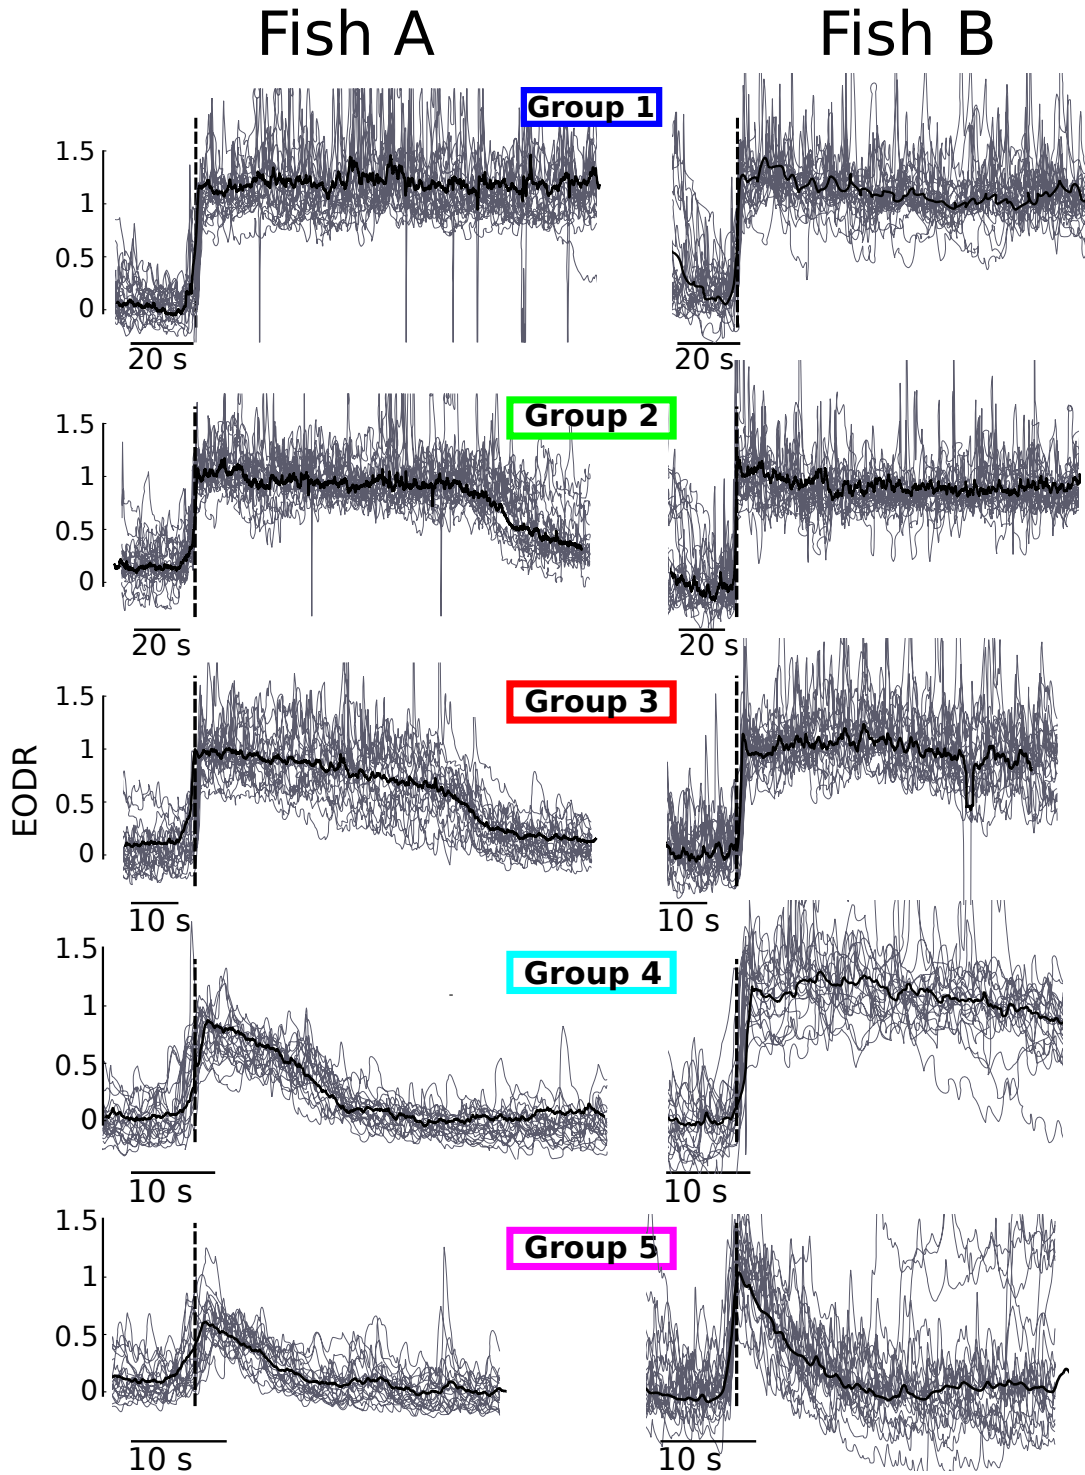


Figure 6: The EODR time series is populated by a heterogeneous set of active state shapes. The left panels show active states that belong to the five groups identified in Figure 5B, all aligned with respect to the movement onset time (dashed vertical lines). The black traces are group averages. The groups of Fish B are extracted in exactly the same way as for Fish A.

filtered time series. Applying the principles of the Schmidt trigger [48], we choose the two halfway points between the local minimum and the local maxima as two distinct thresholds for detecting either movement onset or offset. An upward crossing of the upper threshold is registered as a movement onset transition, and a downward crossing of the lower threshold as a movement offset time.

Active State Features

Here are the details on how we define and calculate the five features used to characterize the shape of active states:

- Transition amplitude, f_1 : Given an inactive-to-active transition, we define two time windows, one after and one prior to the transition, both with a duration of 30 seconds. The transition amplitude is calculated as the difference between the average EODR within these two time windows. Note that the last 5 seconds of the first time window are neglected due to the presence of a short (< 5 seconds) preparatory increase of the EODR prior to the transition.
- State duration, f_2 : Calculated as the difference between the movement offset and movement onset times.
- Average EODR, f_3 : Given by the average of the EODR over the duration of the active state.
- Variance of the EODR, f_4 : Given by the variance of the detrended EODR over the duration of the active state. Detrending (Matlab's 'detrend' function) is necessary since some active states show a slight downward trend.
- Decay duration, f_5 : Given an active-to-inactive transition, the decay duration is calculated as the time taken for the EODR to decay back to its baseline value, as calculated by the same averaging time window used for f_1 , prior to the movement onset.

Once PCA is performed on this feature space, we obtain the following eigenvector for the PC1: for fish A, $PC1 = 0.47\hat{f}_1 + 0.42\hat{f}_2 + 0.40\hat{f}_3 + 0.50\hat{f}_4 + 0.41\hat{f}_5$, for Fish B, $PC1 = 0.28\hat{f}_1 + 0.61\hat{f}_2 + 0.63\hat{f}_3 + 0.36\hat{f}_4 + 0.12\hat{f}_5$, where \hat{f}_i represents the unit vector for the associated axis.

State Segregation and Group Definition

Once the active states have been discretized into the five features and transposed onto principal component space, we use their ordering along the PC1 axis to define groups of neighbouring states that are of similar shape. Group 1 and Group 5 comprise the 20 states with the highest and lowest PC1 value, respectively, while the 20 states that belong to Group 2, 3, and 4 are located above specific relative values of the PC1: 4/5, 3/5, and 1/6 of the maximum PC1 value, respectively. These values are arbitrary and have no impact on the model itself, they merely affect the appearance of the resulting groups, which are only used for visualization purposes. We choose these specific values in order to adequately showcase the full range of different shapes that active states adopt in the data.

2.2.5 Derivation of the Nonlinear Stochastic Model

We outline here the details of the nonlinear stochastic model that we propose to characterize the aforementioned neural control process. Given the lack of experimental evidence available to biophysically constrain this process, it would be premature at this point to develop a detailed neural network model for it. There is rather a preliminary need to characterize the key dynamical features of the data from a phenomenological perspective. To achieve this, we propose a model with the simplest combination of dynamical components that most closely reproduces the statistics of the EODR time series. This model obeys the following stochastic differential equation:

$$\frac{dx}{dt} = -\frac{\partial U(x, t)}{\partial x} + \sqrt{2D} \cdot \Gamma(t) \quad (41)$$

where x is the simulated EODR, D is the noise intensity, $\Gamma(t)$ is Gaussian white noise with mean zero and autocorrelation $\langle \Gamma(t)\Gamma(t') \rangle = \delta(t - t')$, and $U(x, t)$ is a non-stationary double well potential function that can, depending on its asymmetry, give rise to bistability between two stable points separated by an unstable point (Figure 7). It adopts the following form:

$$U(x, t) = a(t)(x - 0.5) + b(t)(x - 0.5)^2 + c(t)(x - 0.5)^4 \quad (42)$$

where an offset of 0.5 is introduced to have the stable points centered on 0 and 1, thereby

matching the format of the EODR. Although all the parameters in Equation (42) are time-dependent, the non-stationarity is in fact mediated by a single underlying variable, $s(t)$, described below.




Figure 7: To account for key aspects of the data, we propose a dynamical system model consisting of a stochastic variable evolving in a bistable potential function that is modulated by a non-stationary, latent variable, $s(t)$. When the tilt parameter, $a(t)$, is zero, the potential function is symmetrical (dashed function). The potential function has two stable points: $x = 1$ represents the active state, and $x = 0$ the inactive state. For Fish A, not only does the potential function tilt back and forth, but the separation variable, $d(t)$, is also modulated. For Fish B, the only source of non-stationarity is the tilt variable, $a(t)$.

The second term in the right-hand side of Equation (41) is responsible for the stochastic nature of the model. Considered on its own, this term would generate a noisy, fluctuating solution, but it would lack any attractor states (i.e. fixed points). This is why we need the first term on the right-hand side of Equation (41), which is the deterministic and nonlinear component of the equation. On its own, and without non-stationary forcing, this term would establish two attractor states, i.e., bistability. This would, however, only yield a trivial solution, namely the decay to the left or right fixed point attractor, depending on whether the starting point was to the left or the right of the origin, respectively. In theory, and with infinite numerical computing accuracy, the solution would converge toward one or the other fixed point indefinitely; in practice, allowing finite precision, one simply says that the noise-free solution has reached either fixed point after a finite transient decay period. Even by adding the non-stationary forcing, the solution would merely become a binary version of the forcing, displaying transitions only when the forcing changes significantly. To display the non-trivial statistics observed in the data,

the model needs the interplay between the three components: stochasticity, nonlinearity, and non-stationarity. In that case, the solution can transition randomly between the two attractor states with a rate that depends on the shape of the potential function at any given time.

The interplay between these three model components is responsible for the existence of a wide range of timescales in the solution of the model: The stochastic component creates very fast (order of 1 second) fluctuations within a single well, adding the nonlinear component allows for switching between wells on the order of 10 to 100 seconds, and the non-stationary forcing slowly modulates the transition rates between wells on the order of 1000s.

When $a(t)=0$, the potential function is symmetrical and the remaining parameters, $b(t)$ and $c(t)$, can be expressed as functions of the shape parameters of this symmetric double well, namely the depth of the wells with respect to the unstable point, h , and the separation, $d(t)$, between the unstable point and either stable points (Figure 7). The relations between these parameter sets are $b(t) = 2h/(d(t))^2$ and $c(t) = -h/(d(t))^4$. When simulating the model numerically, we use these expressions as a practical way to specify the shape of the potential function $U(x, t)$: we first choose a value for h and a time-dependency for $d(t)$, which are then used at every time step to calculate the canonical parameters $b(t)$ and $c(t)$. These values specify the symmetric part of $U(x, t)$. Additionally, a time-dependency for $a(t)$ is prescribed and superimposed on this symmetric function, inducing a tilting of $U(x, t)$ and thereby completing the definition of the deterministic part of Equation (41), for a given time step.

Based on visual inspection of the data, we choose, for Fish A, the tilt parameter, $a(t)$, and the separation, $d(t)$, to be both linearly rescaled versions of a slow latent variable, $s(t)$ (see Eq. (43)). This configuration results in a potential function that is tilted toward the active state ($x = 1$) and has a large separation between the stable points for high values of $s(t)$. Conversely, for low values of $s(t)$, $U(x, t)$ is tilted toward the inactive state ($x = 0$) and has a lower separation between its stable points. Also based on visual inspection of the data, we keep constant the depth, h , while the separation, $d(t)$, is constant for Fish B, but remains time-dependent for Fish A.

We prescribe two different time-dependencies for the slow variable, $s(t)$, depending on which analysis is being carried out (for details on how $s(t)$ is specified in both cases, see Section 2.2.5). In the first case, we run Monte Carlo simulations on a short timescale around active-to-inactive transitions and thus simply prescribe a linear decay for $s(t)$, which we refer to as $s_{local}(t)$. In the second case, simulations are performed over the timescale of the experiments (several hours), which warrants a more realistic $s(t)$. To obtain a variable that could realistically represent the slow driving force that modulates the potential function, we apply a moving average filter with a large window to the EODR time series and we use this filtered trace as $s(t)$. The rationale behind this method is that a smoothed version of the EODR contains the desired information on a slow timescale and thus constitutes a first estimate for the latent variable, should it exist.

Although this moving average filter will smooth out abrupt transitions, the simulations driven by this latent variable will still show transitions as abrupt as in the data, i.e., on a much faster time scale than the one on which the latent variable fluctuates. The slower rise or decay of the latent variable will, however, allow transitions to occur at a rate that is slowly modulated by this latent variable, as is observed, for instance, in the data following long active states.

Also note that by using the experimental data to infer the latent variable, and then using this latent variable to drive the model, we inevitably impart some level of correspondance between the data and the model results. This will, however, only be the case for timescales longer than the averaging window used to estimate the latent variable (~ 2300 seconds for Fish A, ~ 340 seconds for Fish B). Hence, out of the three timescales mentioned above, only the first two emerge directly from the model. For instance, the longest active states observed in the data (such as those shown by the black arrows in Figure 8A) will automatically be reproduced by the model, since they yield a very high value for $s(t)$ when averaged out, which in turn allows a single attractor (namely, the one to the right of the origin) to exist in the model.

Our modelling framework and our interpretations are predicated on the following set of assumptions:

1. The spontaneous decision to initiate or terminate movement emerges from an open-loop subsystem, i.e., it is intrinsically generated by a high-level neural population.




Figure 8: Simulation results from Model A (lower panel, blue trace) qualitatively match the data from Fish A (upper panel, blue trace). This segment of data is taken from the trace in Figure 3. Applying a moving average filter to the observed EODR yields the black trace, which is then used as the latent variable, $s(t)$, for Model A (see Section 2.2.5). The black arrows are examples of active states that are expected to be reproduced by the model since they are longer than the averaging window used to obtain the $s(t)$.

2. The activity of this subsystem can be projected onto a low-dimensional manifold containing bistable attractor dynamics.
3. Neural noise causes this subsystem to randomly alternate between the two stable modes of activity, or attractor states.
4. Modulatory agents, e.g., monoamines or peptides, evolving on a slow timescale, tilt the bistable attractor landscape back and forth and thus affect the animal's propensity to move.
5. Information on the activity of this subsystem is conveyed through the EODR. The EODR thus represents a proxy for the neural activity responsible for triggering spontaneous transitions between the active and inactive behavioural states.

Assumption 1 is plausible given that the experiments were carried out without any external cues, precluding any sensory-evoked responses. The plausibility of Assumptions 2 and 3 can be established in light of the fact that i) recurrent networks have been shown to generate bi- and multistable attractor dynamics, in a winner-take-all fashion [198, 126] ii) it has been suggested that neural noise can be responsible for triggering transitions between the attractor states [126], and iii) a bistable attractor network model [198] can

be reduced to a one-dimensional stochastic differential equation of the same form as Equation (41) [167]. Furthermore, experiments on zebrafish have identified a clear link between neuropeptidergic modulation and arousal behaviour [157, 204], which supports Assumption 4. Assumption 5 was discussed in the Introduction and is consistent with previous findings [203, 60, 95].

Integration of the Stochastic Differential Equation

To integrate Equation (41), the asymmetry parameter, $a(t)$, the separation, $d(t)$ (or $d = \text{constant}$ for Fish B), the height h , and the noise intensity, D , need to be specified. Both $a(t)$ and $d(t)$ are linearly rescaled versions of $s(t)$:

$$\begin{aligned} a(t) &= \frac{(a_2 - a_1)}{(s_{max} - s_{min})} \cdot (s(t) - s_{min}) + a_1 \\ d(t) &= \frac{(d_2 - d_1)}{(s_{max} - s_{min})} \cdot (s(t) - s_{min}) + d_1 \end{aligned} \quad (43)$$

where s_{max} , s_{min} , a_1 , a_2 , d_1 , d_2 are the extremum values of $s(t)$, $a(t)$, and $d(t)$, respectively. Note that s_{max} and s_{min} are not free parameters, but are rather obtained from the $s(t)$ time series. Parameter values for both fish are shown in Table 4. Parameters are obtained sequentially under various constraints that minimize the differences between simulations and experimental results. First, d_1 and d_2 are chosen to match the minimum and maximum transition amplitude calculated from the data (Figure 9, Group 5 and Group 1, respectively). Note that assigning $d_1 = d_2$ for Fish B yields a constant $d(t)$. Once the separation is fixed, h can be used to adjust the slope of the potential function between the stable and unstable point, and therefore controls the abruptness of the transitions between stable points. It is chosen such that simulated transitions are as abrupt as those observed in the data. Transitions from Fish B are slightly steeper than those of Fish A (Figure 9), which is why we obtain a larger depth for the simulations of Fish B. a_1 and a_2 are chosen such that the potential function, in its most asymmetric configurations, allows states and transitions that are similar to those of the data. For instance, when $s(t)$ adopts its highest values, the experimental data for both fish show active states that are long, with rare instances of brief inactive states (Figure 8 and Figure 10, top panels). This observation constrains the value of a_2 : when $a(t) = a_2$, the potential function must have a deep active state well along with a shallow inactive state well. On the other

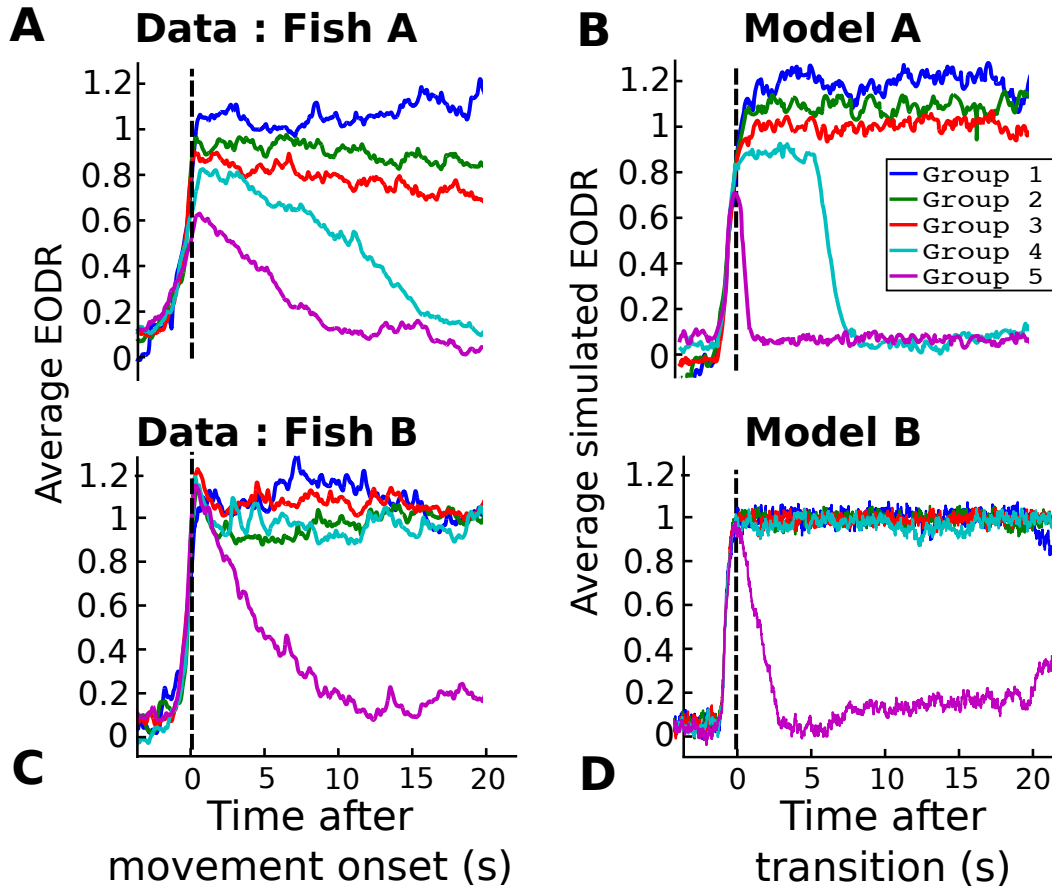


Figure 9: The models for both fish capture the essential features of inactive-to-active transitions across different groups. Coloured traces represent the average of a given group, i.e., the traces from panel A and C correspond exactly to the black traces of Figure 6. All active states are aligned with respect to movement onset time, in the case of experimental data, or with respect to the time of an upward crossing of the model’s unstable point, in the case of simulation (vertical dashed lines).

hand, when $s(t)$ adopts its lower values, this situation is reversed in the case of Fish A, while for Fish B, the EODR alternates between active and inactive states, with a bias for the inactive state (Figure 10, upper panel). This argues for a potential function that is slightly more symmetrical than that of Fish A, which is why the value of a_1 is lower for Fish B (recall that $a(t) = 0$ yields a symmetrical potential function). The remaining free parameter, D , controls the switching rate between the stable points. It is therefore obtained such that the number of transitions in the simulations matches that seen in the experimental data over the same period of time, and with the same $s(t)$. Numerical integration is performed with the Euler–Maruyama scheme with a time step of 0.01 seconds.




Figure 10: Simulation results from Model B (lower panel, blue trace) qualitatively match the data from Fish B (upper panel, blue trace). Applying a moving average filter to the observed EODR yields the black trace, which is then used as the latent variable for Model B (see Section 2.2.5).

Table 1: Numerical integration parameters for the models of both fish

Parameter	Fish A	Fish B
h	-0.08	-0.32
d_1	0.3	0.5
d_2	0.6	0.5
a_1	0.12	0.07
a_2	-0.1	-0.1
D	0.02	0.1

Estimation of the Latent Variable

The latent variable, $s(t)$, is responsible for the non-stationarity imparted to the model, i.e. a constant $s(t)$ would turn Equation (41) into a stationary stochastic process. Depending on which analysis is considered, we prescribe two different forms for $s(t)$: it is either a slow and stochastic time series evolving over a long timescale, or a linear decay over a short timescale. The latter is referred to as $s_{local}(t)$.

For the longer simulations performed in Sections 2.3.1 and 2.3.3, the stochastic differential equation (41) is integrated over a long period of time where the latent vari-

able, $s(t)$, is chosen as the moving average filtered version of the EODR time series (Figure 8 and Figure 10, black traces). The window size for the filter is chosen as $4(\langle T_{act.} \rangle + \langle T_{inact.} \rangle)$, where $\langle T_{act.} \rangle$ and $\langle T_{inact.} \rangle$ are the mean residence times of the active and inactive state, respectively. The window size for Fish A is 2997 seconds, while for Fish B it is 344 seconds. The window size is chosen large enough such that a change in the filtered EODR implies that the durations of the states undergo a significant, sustained deviation from their average values, which relates to the level of tilt of the potential function. The window size is also chosen such that $s(t)$ becomes approximately constant when the true latent variable is constant, i.e., when the process is stationary (as tested with simulations).

The Monte Carlo simulations performed in Section 2.3.2, on the other hand, are performed over only 1500 seconds, with a prescribed linear decay for $s_{local}(t)$. This represents a potential function, tilted toward the active state at first, that undergoes a tilting toward the inactive state at a constant rate. This rate, i.e. the slope of the linear decay of $s_{local}(t)$, is chosen as -0.036. This value is obtained under the constraint that the transient period of bistability in simulations is as long as that seen in the data (Figure 12). A higher tilting rate (i.e., a steeper slope for $s_{local}(t)$) would shorten this bistable period, whereas a lower rate would lengthen it. Since the influence of $s_{local}(t)$ is mediated only through linearly rescaled versions of itself, i.e., $a(t)$ and $d(t)$, the y-intercept of the linear profile is arbitrary.

2.3 Results

Below we present three distinct modelling experiments aimed at validating our proposed modelling framework. With the first two experiments, we examine trajectories of the EODR around the transitions between states. The third experiment shows how our simulations yield residence time distributions consistent with those seen in the data.

2.3.1 Onset-triggered Analysis

Here we examine how the EODR behaves around the transitions from the inactive to active state, that is, around the times of movement onset. We first investigate the correlations between the transition amplitude (f_1) and the active state duration (f_2) by

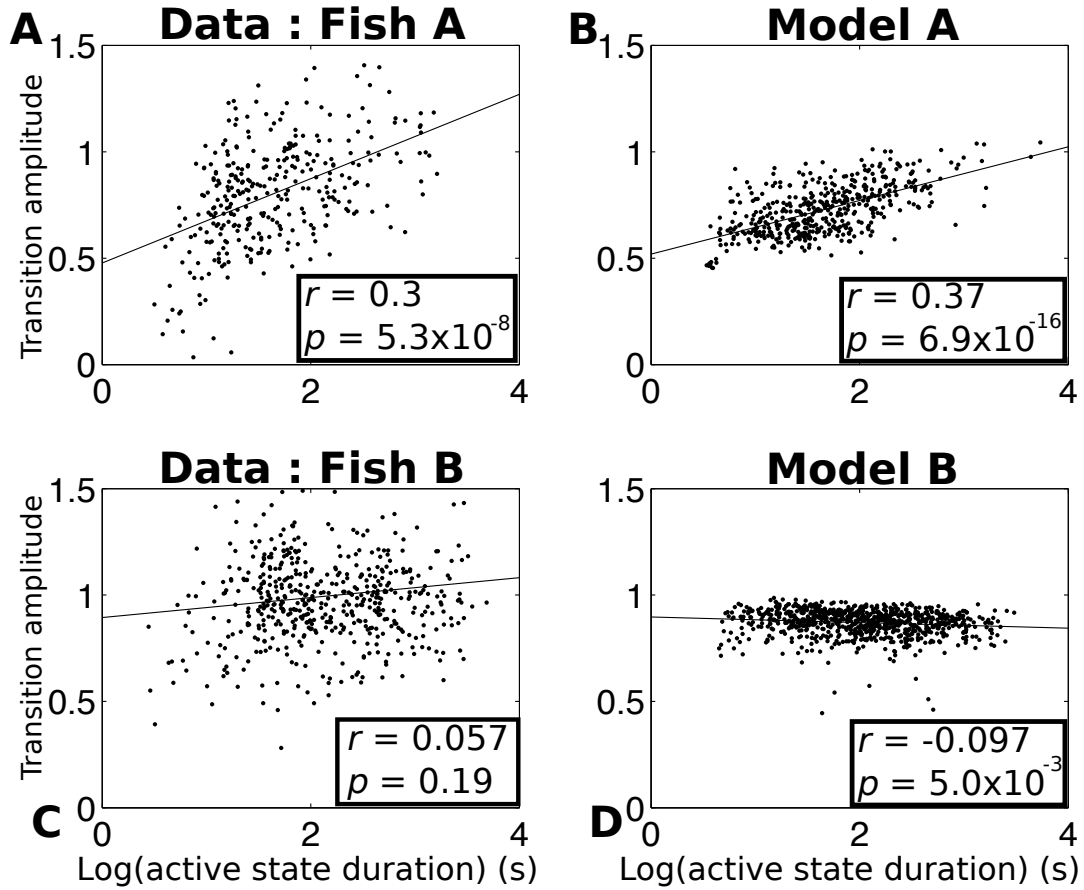


Figure 11: The model captures the correlation (Fish A), or lack thereof (Fish B), between the transition amplitude (f_1 , y-axes) and the duration of the active state (f_2 , x-axes). Each black dot represents an active state, projected on the $f_1 - f_2$ plane. For visual reference, a least-square line is also shown (straight black lines). The linear correlation coefficient and the p-value for each plot are also reported as r and p , respectively.

projecting the entire feature cluster of Figure 5A onto the $f_1 - f_2$ plane, such as to generate 2-D scatter plots. Comparing the scatter plots for both fish (Figure 11A and Figure 11C) reveals that, for Fish A, the transition amplitude covaries with the active state duration across the data set, with shorter states having the lowest amplitude and longer states the greatest. For Fish B, however, the transitions from inactive to active states have an amplitude that is uncorrelated with the state duration.

The unexpected variability in the dynamics of inactive to active state transitions implies important differences in neural network dynamics across individuals of the same species. This discovery invalidates blind across-individual averaging of behavioural or neural data and, as described below, requires individual-specific model choices.

To reproduce these results with our stochastic model, we first integrate Equation (41) for a duration similar to that of the pooled EODR time series, ~ 60 hours. In this case, the latent variable, $s(t)$, is obtained by applying a moving average filter to this pooled time series (see Section 2.2.5). This method yields a realistic $s(t)$ that evolves stochastically over a timescale much slower than that of the actual EODR (Figure 8 and Figure 10, black traces).

For Fish A, in addition to a time-dependent tilt parameter, $a(t)$, we also assign a time-dependency to the separation, $d(t)$, both of which are linearly rescaled versions of the latent variable, $s(t)$. This configuration creates a greater separation between the two stable points when the potential function is tilted toward the active state, and a smaller separation when it is tilted toward the inactive state. Note that, by construction, the separation, $d(t)$, directly controls the transition amplitude. Also, since the level of tilt of the potential function strongly controls the durations of the active states (with, e.g., longer active states occurring when the potential function is tilted toward $x = 1$), simulations will show a positive correlation between the transition amplitude and the active state duration because $a(t)$ and $d(t)$ are both dependent on the same variable. For Fish B, on the other hand, only $a(t)$ is dependent on $s(t)$, with $d(t)$ held constant.

These choices are based on the observations, described in the previous paragraph, that the transition amplitudes correlate with the active state durations for Fish A, but not for Fish B. It should be noted that, although the shape parameters, $d(t)$ and h , of the (symmetric) double well can be held constant, the presence of a varying tilt parameter, $a(t)$, inevitably causes fluctuations in the actual well depth and well separation of the potential function. We found, however, that these fluctuations are not by themselves able to explain the range of transition amplitudes observed in Fish A, which is why a time-dependent separation, $d(t)$, was introduced for that case.

The time series obtained from these simulations qualitatively match those of the observed EODR (Figure 8 and Figure 10) and are processed in exactly the same way as the experimental data, i.e. they are fed through the state segregation scheme described in Section 2.2.4. In the case of simulations, however, the EODR variance in the active state, f_4 , and the decay duration, f_5 , remain constant because variability of these features are

not included in the model. Nevertheless, an analysis of these simulations yields scatter plots for both models that capture the correlation, or lack thereof, between the transition amplitude and the active state duration (Figure 11B and Figure 11D). By comparing model results and data in Figure 11, it is apparent that the observed transition amplitude has a greater variability than the simulated one. This is likely caused by the low number of parameters of the potential function used in our model, which makes it tighter than the “true” potential function. Adding parameters that could widen the bottom of the wells would allow for more variability in the simulated transition amplitudes.

By using the segregation into groups described in Section 2.2.4, we can plot the average traces for each group, observed and simulated (Figure 9). We arrive at the same conclusion drawn from Figure 11, i.e., that Fish A transitions from inactive to active states with an amplitude that covaries with other attributes of the active states, as shown by the fact that different groups have different average transition amplitudes. In contrast, for Fish B, the transition amplitude is independent of the group. In addition, Figure 9 shows that the simulated time series exhibit a similar spectrum of active state shapes as the data.

From Figure 6 and Figure 9, one sees that some active states from the experimental data tend to follow a slow, almost linear decay toward the inactive state, which contrasts with the steeper upward transitions toward active states, see e.g., Group 4 of Fish A in Figure 6. The analysis presented in this section does not attempt to reproduce this asymmetry, which is why the active states from the simulations show active-to-inactive transitions that are more abrupt than those of the data, as seen from the average traces of Group 4 and Group 5, Figure 9. See Section 2.4 for potential additions to the model that could capture this asymmetry. In the next section, however, we show that a separate analysis, centered specifically around the movement offset times, offers a satisfying fit between simulated and observed active-to-inactive transitions.

2.3.2 Offset-triggered Analysis

We now examine in detail the structure of the EODR around movement offset times, that is, around active-to-inactive transition, for active states that belong to Group 1, Fish A. This particular group of active states warrants special attention because it possesses

attributes that convey unique information about the mechanisms that could govern these transitions.

For this particular group, we first wish to obtain the non-stationary probability density function (PDF) of the EODR, time-locked to the time of movement offset, which we refer to as the transition-triggered PDF. This is obtained by first aligning each active state of this group with respect to the time of movement offset, and chopping off data that fall outside a 300s time window centered around that time (Figure 12, upper panel). The remaining set of EODR traces is then divided into several time bins, and within each bin we obtain a PDF for this segment of the EODR traces. Putting all these segments together yields a time-dependent PDF for the EODR, distributed around the time of movement offset (Figure 12, lower panel). This analysis reveals two defining features of active-to-inactive transitions for this group of active states: i) on average, the EODR undergoes a downward trend starting ~ 1000 s before movement offset, and extending ~ 1500 s beyond it (although in Figure 12 we only show ~ 500 seconds around movement offset); ii) there is a significant probability that the fish briefly returns to an active state almost immediately after movement offset. These traits are absent from the other groups of active states and from other fish.

To explain these results, we propose a dynamical scenario that involves a transient period of bistability caused by a slowly tilting double well potential. The associated deterministic dynamical system starts off with a single stable state that corresponds to the active state, then undergoes two saddle-node bifurcations that cause the appearance of the inactive stable state followed by the disappearance of the active stable state, and finally adopts its final configuration in the inactive state (Figure 13, upper panel, red and green traces).

To show that this explanation reproduces the experimental results described above, we perform Monte Carlo simulations of Equation (41) under this scenario. In this case, because the timescale is so short, we simply prescribe a linear decay for the latent variable, $s_{local}(t)$, which yields a constant tilting rate for the double well potential function. This rate is varied until the simulations most closely match the experimental data (see Section 2.2.5). The system is initialized in the active state and forced by this progressive tilting. To complement the Monte Carlo approach, we also solve the Fokker-Planck equa-

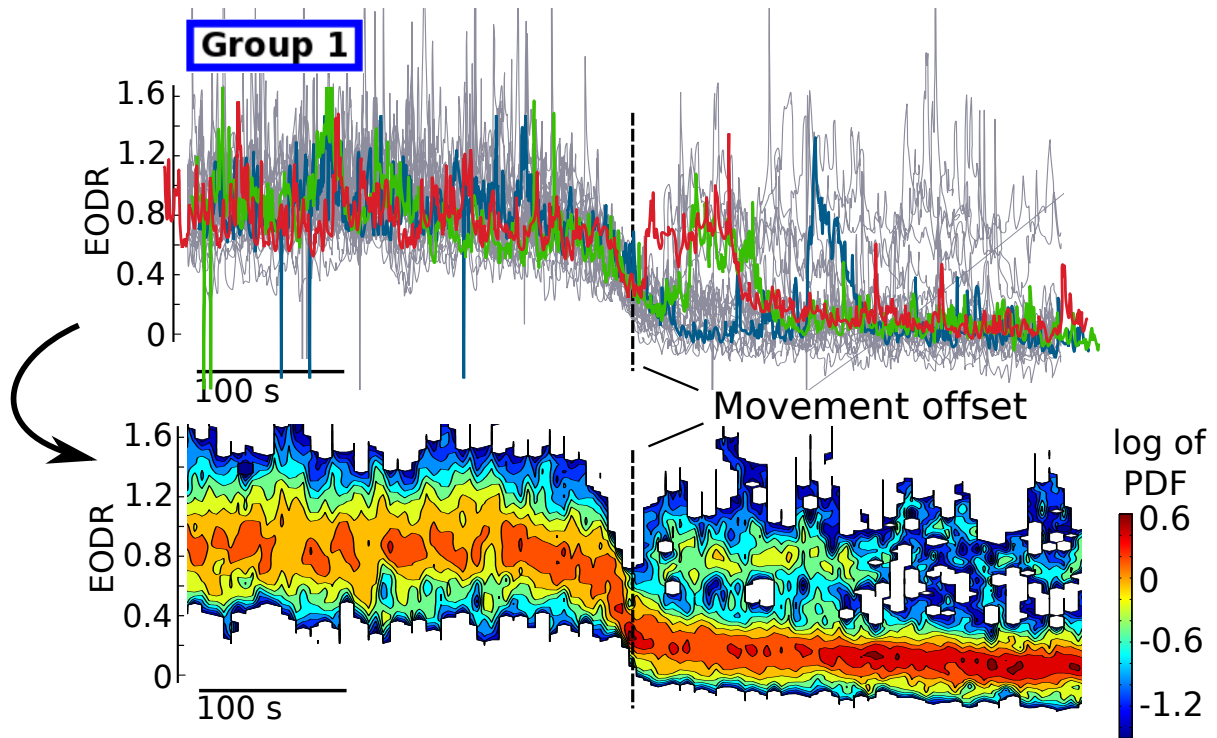


Figure 12: Group 1, from Fish A, shows a high propensity to return briefly to the active state shortly after the time of movement offset. This is superimposed on a slow, tail-like decay of the EODR. The grey traces in the upper panel correspond to the same active states that are shown in the upper left panel of Figure 6, but in this case, they are aligned with respect to the movement offset time, which allows a transition-triggered PDF to be compiled around movement offset (lower panel). Three representative traces in the upper panel are shown in colour as examples.

tion associated with this non-stationary stochastic process and those initial conditions. This solution confirms the presence of a brief, transient bistability period intercalated between two monostable regimes (Figure 14).

For each realization of the non-stationary stochastic process, we assign a timestamp to the moment when the unstable fixed point of the system is crossed, and then align all the Monte Carlo realizations of the process with respect to this time (Figure 13). The same procedure was applied to the experimental data in Figure 12, except that the movement offset time was used as the reference. Note that the realisations found in the upper panel of Figure 13 follow the non-stationary PDF of Figure 14. Once the realisations are aligned with respect to the transition times, however, they can no longer be compared to Figure 14, because they are now conditioned on a transition at time 0, rather than on the initial conditions used to solve the Fokker-Planck equation.




Figure 13: Top: Set of 30 Monte Carlo simulations of Equation (41), driven by a linear decay of the latent variable, $s_{local}(t)$. The stable and unstable point of the system are shown in green and red, respectively, and the black stars are time stamps marking downward crossings of the unstable point. The system is initialized at $x=1$ and allowed to stabilize before the potential landscape starts to tilt over. All the realisations in the upper panel obey the PDF shown in Figure 14. Bottom: All traces from the upper panel are aligned with respect to the time stamps, which then allows the transition-triggered PDF of Figure 15 to be obtained.




Figure 14: The non-stationary solution of the Fokker-Planck equation associated with Equation (41), driven by a linear decay of the latent variable, $s_{local}(t)$. The solution confirms the presence of a transient bistability period (dashed rectangle) that allows the system to briefly jump back to the active state immediately after a downward transition. The slope of $s_{local}(t)$, i.e., the tilting rate, controls the duration of this bistability period, with a higher rate leading to a briefer period. This solution is obtained numerically with a custom partial differential equation solver using finite volume discretization and implicit time-stepping. Traces from the upper panel of Figure 10 evolve according to this PDF.

With the simulated EODR traces properly aligned, we now complete the analysis by obtaining the transition-triggered PDF for the simulated EODR (Figure 15). This distribution is obtained in exactly the same way as the one shown in Figure 12 for the experimental data. The resulting comparison between Figure 15 (simulations) and Figure 12 (data) is satisfactory, with the model correctly capturing the slow downward trend of the EODR, as well as the tendency to return to the active state immediately after the transition. Note that by fine-tuning parameters associated with these Monte Carlo simulations, the match between data and model could be improved. Notably, by prescribing a non-linear decay for $s_{local}(t)$, comprising, e.g., a plateau for the second half of the simulations, some bistability could remain for a longer time, allowing active-to-inactive transitions to occur more than 100 seconds following the transition time, similar to what is seen in Figure 12.




Figure 15: Monte Carlo simulations of the transiently bistable system yields a transition-triggered PDF that is qualitatively similar to that obtained from the data. The 30 iterations used to generate this distribution are first aligned with respect to the time when the unstable point of the system is crossed downward, and are then processed in the same way as the experimental data to obtain the transition-triggered PDF shown here.

Transition-Triggered PDF

With the active states from a given group aligned with respect to either the downward crossing of the unstable point (simulations), or the movement offset time (experimental

data), we seek a PDF for the EODR, or simulated EODR, conditioned on a transition at a given time: for simulations, we seek $p(x, t | \text{'unstable point crossing at } t=0\text{'})$, and for experimental data, $p(x, t | \text{'movement offset at } t=0\text{'})$, where x is the simulated or observed EODR, $t \in [-T/2, T/2]$, and where we choose $T \sim 500$ seconds. To obtain this density, we bin each individual traces in 5 second windows, such that with our temporal resolution of 0.01 second and with 20 traces per group, we have access to 10,000 values per time bin. For each bin, we then obtain an estimate for the PDF by using Matlab's kernel density estimator ('ksdensity' function, default options). Putting all time bins together, we can show the evolution of the PDF with a colorplot, centered on the desired transition time.

Numerical Integration of the Fokker-Planck Equation

The Fokker-Planck equation associated with Equation (41) is

$$\frac{\partial p(x, t)}{\partial t} = \frac{\partial}{\partial x} \left[\frac{\partial U(x, t)}{\partial x} p(x, t) \right] + D \frac{\partial^2 p(x, t)}{\partial x^2} \quad (44)$$

where $p(x, t)$ is the PDF of $x(t)$ given the initial condition $x(t = 0) = 1$, and where the non-stationarity embedded in $U(x, t)$ is conveyed by the latent variable, $s(t)$. We obtain the non-stationary solution of this equation by numerical integration with a custom partial differential equation solver that implements a finite volume discretization with the fully implicit Euler scheme. The advective term is treated with the upwind scheme and a linear interpolation profile for the spatial derivative of $p(x, t)$ is applied to the diffusive term [150]. The resulting algebraic equation is solved with the tridiagonal matrix algorithm [150]. The spatial resolution Δx is 0.004, for a total of 1000 grid points between $x = -2$ and $x = 2$. The time step is the same as for the integration of Equation (41), 0.01 second. Note that, although the solution of this equation (Figure 14) looks similar to the transition-triggered PDF obtained from the simulations (Figure 15), they should not be compared together. Although they are both non-stationary PDFs, they are not conditioned on the same event: the former is conditioned on $x(t = 0) = 1$, with $t \in [0, T]$, while the latter is conditioned on a transition occurring at $t = 0$, with $t \in [-T/2, T/2]$.

2.3.3 Residence Time Distributions

For this final analysis, we focus on the residence time distributions (RTD), that is, the PDF for the time spent in a given state. We show that our proposed modelling framework produces RTDs that are consistent with those of the data. Here we focus primarily on the inactive state RTDs. This is because fish movement during the active state might cause re-afferent signals that would interfere with the neural process governing the termination of movement. Notably, the possibility of the fish encountering the tank walls is an unavoidable element that might invalidate our assumption that an open-loop subsystem is responsible for terminating movement (see Assumption 1 in Section 2.3). During the inactive state, however, the fish’s sensing volume remains unchanged and our working assumptions hold. Unless specified otherwise, the inactive state is thus implicitly assumed for the remainder of this section.

Recall that the experimental data from each fish is divided into several recording sessions. Comparing the RTDs obtained from each of these sessions with the RTD of the pooled time series reveals significant statistical differences between them. This prevents us from assuming that the residence times from different recording sessions are sampled from the same underlying PDF. We thus refrain from pooling those residence times and rather opt for analysing each recording session separately.

For all fish, we find that the stretched exponential family of PDFs provide a satisfying fit for the RTDs (Table 2, second column). Stretched exponential distributions lie on a continuum between exponential and power-law distributions, with a single parameter controlling which regime is more expressed. Except for the limiting exponential case, they possess a long tail and are often used to describe scale-free phenomena [122].

Table 2: Comparison between the fitting statistics for the inactive state residence time distributions of the data and of model results, for each recording session.

Data					Simulations				
Fish-Session	p-val.	# of states	α	$\langle t \rangle$ (s)	p-val.	# of states	α	$\langle t \rangle$ (s)	
A-1	0.27	49	0.30	473.14	0.31 ± 0.34	43.74 ± 6.23	0.41 ± 0.34	473.21 ± 81.43	
A-2	0.37	30	0.27	523.30	0.31 ± 0.32	32.34 ± 4.49	0.33 ± 0.22	466.15 ± 77.94	
A-3	0.03	51	0.24	773.87	0.25 ± 0.31	68.21 ± 7.47	0.34 ± 0.20	508.46 ± 64.71	
A-4	0.59	93	0.29	529.59	0.20 ± 0.28	80.59 ± 8.66	0.38 ± 0.21	542.10 ± 75.30	
A-5	0.45	67	0.47	609.23	0.24 ± 0.30	67.76 ± 7.01	0.38 ± 0.27	542.22 ± 66.51	
A-6	0.67	31	0.52	524.04	0.27 ± 0.31	36.82 ± 5.24	0.44 ± 0.37	414.97 ± 68.59	
A-7	0.71	72	0.42	439.88	0.17 ± 0.25	72.38 ± 8.14	0.32 ± 0.20	408.15 ± 60.43	
B-1	0.65	180	1.00	32.18	0.47 ± 0.26	73.70 ± 6.62	0.56 ± 0.17	53.33 ± 7.42	
B-2	0.77	99	0.93	25.05	0.44 ± 0.25	55.30 ± 7.72	0.62 ± 0.24	42.12 ± 7.21	
B-3	0.51	81	0.77	34.65	0.54 ± 0.26	71.53 ± 6.52	0.71 ± 0.22	40.26 ± 5.23	
B-4	0.78	109	0.61	34.36	0.49 ± 0.25	56.20 ± 6.22	0.57 ± 0.19	51.83 ± 8.30	
B-5	0.37	35	0.55	34.55	0.44 ± 0.25	36.89 ± 5.89	0.58 ± 0.30	40.94 ± 9.41	
B-6	0.91	72	0.84	77.18	0.49 ± 0.25	76.27 ± 7.75	0.57 ± 0.15	73.89 ± 10.15	
B-7	0.69	77	0.45	50.97	0.45 ± 0.24	70.58 ± 7.70	0.54 ± 0.16	58.93 ± 8.74	
B-8	0.75	34	0.33	71.89	0.47 ± 0.27	61.93 ± 7.36	0.55 ± 0.18	63.61 ± 10.50	

The p-values are from two sample Kolmogorov-Smirnov tests between the empirical RTD and the best fit stretched exponential distribution. α and $\langle t \rangle$ are parameters of the stretched exponential distribution of Equation (46). Simulation results are formatted as 'mean \pm std. dev.', as obtained from 100 iterations of Equation (41), driven by the $s(t)$ associated with each recording session.

To determine if our modelling framework also fits RTDs with stretched exponential distributions, we use the same approach as in the Onset-triggered Analysis, where the latent variable $s(t)$ is given by the smoothed EODR (Figure 8 and Figure 10, black traces). The only difference in this case is that we obtain several traces for $s(t)$, one for each recording sessions, rather than extracting a single trace from the pooled EODR time series. Each $s(t)$ trace is then used to drive Equation (41) and thus to obtain a simulated EODR associated with a given recording session. These simulation results are then piped through the same transition detection algorithm as the data, from which we extract the residence times and compile the RTDs for each recording session. The resulting RTDs are then fitted to stretched exponential distributions by the same procedure applied to the data. To obtain proper statistics on the simulated RTDs, we repeat this procedure to obtain 100 iterations of the simulated EODR. From this ensemble of simulation results, we extract, for each recording session, the average p-value evaluating the stretched exponential fitting, the average number of states, and the average parameters of the fitting distribution. We report these average values in Table 2, and compare them to those of the experimental data associated with the same recording session.

Not only are the simulated RTDs well-fitted by stretched exponentials like the data, but they also closely resemble the RTDs of their associated recording sessions, regardless of which fish is considered, as shown in Figure 16A for two examples of recording sessions. As mentioned in Section 2.2.5, however, the fit between the simulated and the data RTDs for large timescales should not be interpreted as model validation, since we have used a large averaging time window to infer $s(t)$ from the data. The length of this time window is shown as the black dotted vertical line in Figure 16. From the ensemble of 100 simulation results associated with each recording session, we can also obtain PDFs for the fitting parameters, α and $\langle t \rangle$, which are then compared to the observed values of these parameters, confirming the correspondance between experimental data and model results (Figure 16B). In most cases (21 parameters out of 30), the observed parameter values fall within one standard deviation of the simulated ensemble average (Table 2). The match between observed and simulated RTDs also holds for the active state, but in this case about half of the recording sessions are not well-fitted by stretched exponentials, although they still possess a long tail.

Lastly, as the appearance of stretched exponential distributions can be associated

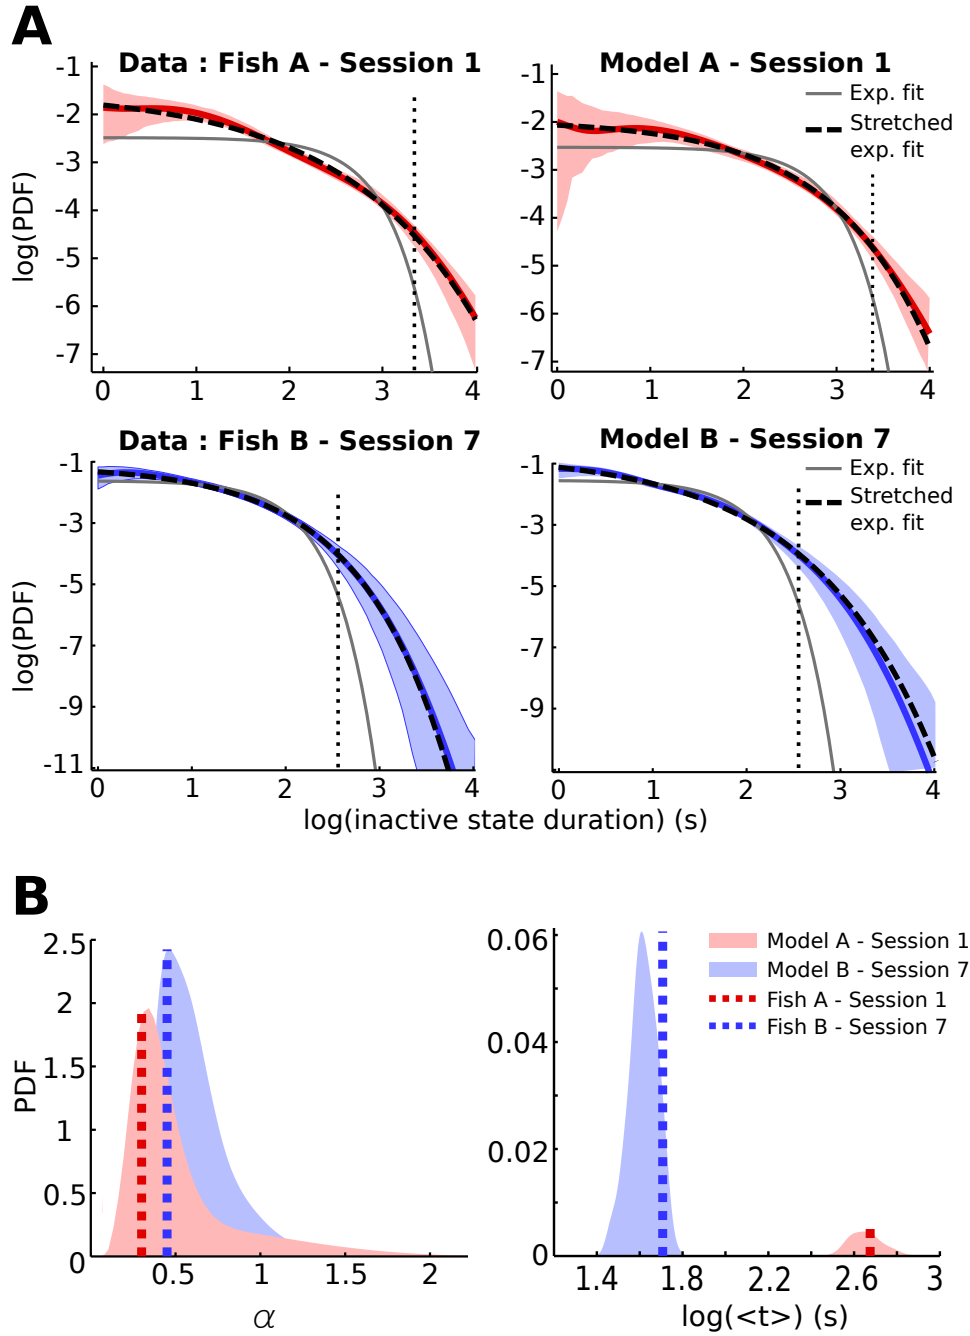


Figure 16: **A)** Inactive state residence time distributions. Models for both fish produce inactive state RTDs (coloured curves) that are consistent with those of the data, all of which are fitted by the stretched exponential distribution functions of Equation (46) (dashed curves). For comparison, gray curves show the best fit exponential densities for the RTDs. Fitting parameters for all sessions can be found in Table 2, for both experimental data and simulations. Shaded areas are 95% bootstrap confidence interval for the RTDs, obtained with Matlab's 'bootci' function (1000 bootstrap samples). The simulated RTDs shown here are obtained from a single realisation of Equation (41), driven by the $s(t)$ associated with the appropriate recording session. The dotted vertical lines correspond to the value of the averaging window used to obtain $s(t)$. **B)** Distribution of fitting parameters. These distributions are obtained from the ensemble of 100 iterations associated with the same sessions as in A. Dashed lines show the value of these parameters obtained from fitting the experimentally observed RTDs. These values correspond to the 4th and 5th columns of Table 2.

with long-range correlations [122], we examine whether there exist correlations between the duration of successive inactive states (i.e., duration of the i^{th} inactive state against that of the $(i - 1)^{\text{th}}$ one). For Fish A, we find no such correlations, in neither model nor data. For Fish B, however, a small but significant correlation is observed (data: $r = 0.13$, $p = 6.69 \times 10^{-4}$; model: $r = 0.23$, $p = 1.89 \times 10^{-15}$). For both data and model these correlations persist up to a lag of 6 to 8 inactive states. This shows that, for Fish B, the latent variable $s(t)$ imparts some degree of memory on a slow time scale to the faster dynamics of the EODR. The fact that these correlations are absent for Fish A can be explained by the lower level of noise, D , for this individual. Indeed, less noise means that less transitions are triggered, which in this case means that the EODR has less opportunities to “sample” and act as a readout for the latent variable. This is in line with the notion of aperiodic stochastic resonance [30], where there is an optimal level of noise at which the fast variable of slowly and aperiodically driven bistable system is an optimal readout of the driving force.

Stretched Exponential Fitting

The stretched exponential distribution is given by:

$$p(t) = K \exp(-\beta t^\alpha) \quad (45)$$

Following Luevano [122], the prefactor K and constant β can be expressed in terms of the exponent α and the mean value $\langle t \rangle$:

$$p(t) = \frac{\alpha b}{\Gamma(1/\alpha) \langle t \rangle} \exp(-(bt / \langle t \rangle)^\alpha) \quad (46)$$

where $b = \frac{\Gamma(2/\alpha)}{\Gamma(1/\alpha)}$. With knowledge of $\langle t \rangle$, the stretched exponential is thus characterized by a single parameter, α . This property is particularly useful for fitting purposes, since we have access to an estimate of $\langle t \rangle$. Given a set of residence times, from either the data or simulations, its empirical PDF, $\tilde{p}(t)$, is estimated via kernel density estimation (‘ksdensity’ function in Matlab, ‘support’ option set to ‘positive’) and the associated mean is calculated as $\langle t \rangle = \int_0^\infty t \tilde{p}(t) dt$. A 1-D grid search along possible values of α is then performed to minimize the squared difference between the estimated density $\tilde{p}(t)$ and the candidate distribution of Equation (46). The quality of the fit is evaluated with

a two sample Kolmogorov-Smirnov test ('kstest2' function in Matlab, default options), whose p-values are reported in Table 2.

2.4 Discussion

We have shown that a low-dimensional modelling framework, consisting of only a small set of dynamical components, captures the core aspects of spontaneous transitions between active and inactive behavioural states in electric fish. Given that the Jun et al. [95] dataset shows signs of non-stationarity, we developed a scheme that segregates similar segments of data together, allowing us to properly investigate the mechanism from which stereotyped shapes of active states emerge. We propose that this data analysis scheme is a useful tool to break down and understand non-stationarity in stochastic data. By applying this scheme, we examined in detail the average structure of the EODR around times of movement onset and offset.

For transitions around movement onset, our analysis revealed a positive correlation between the transition amplitude and the duration of the active states (for Fish A). By introducing a time-dependency for both the tilt and separation variable of the potential function, our model correctly captures the observed correlation. In the case of transitions around movement offset, our data analysis scheme allowed us to observe the presence of a brief period, immediately following movement offset, where fish show a propensity to return to the active state. To explain this, we proposed a simple dynamical scenario where the potential function is tilted toward the inactive state with a constant rate. This creates a brief period of bistability that allows the simulated EODR to briefly return to the active state following an active-to-inactive transition, as seen in the data. Finally, we showed how simulating the model over long time scale yields time series where the RTDs are long-tailed and well fitted by stretched exponential distributions, and where correlations emerge between the duration of successive states (for Fish B only). Both these modelling results are consistent with the data.

Together, these analyses lead us to conclude that low-dimensional, bistable neural dynamics underly the emergence of spontaneous transitions between behavioural active and inactive states in electric fish, and that stochastic threshold crossing is the mechanism triggering these transitions. These findings corroborate recent results that also

confirm a key role for fluctuating neural activity in the timing of walk/rest transitions in freely-walking *Drosophila* [123]. In addition, a major conclusion of our work is the identification of a non-stationary latent variable that exerts its influence through the shape of the bistable potential landscape governing the transition dynamics. We hypothesize that this bistability is established in highly recurrent telencephalic networks, and that neuromodulation deriving from diencephalic peptidergic systems provides the inferred non-stationary forcing to this network [192, 43].

2.4.1 Putative Neural Structures Involved in the Transitions

Lesions of the telencephalon leave electric fish in the inactive state [152], suggesting that it contains the circuitry responsible for the stochastic switching between inactive and active states. In support of this hypothesis, a recent study has shown that the recurrent networks of the *gymnotiform* telencephalon can be induced to switch between Up and Down states [43]. Furthermore, work in zebrafish suggests that the orexin and galanin peptidergic systems of the hypothalamus can control the probability of such switches [157, 204]. The slow fluctuations of the latent variable that modulates the relative depth of the wells in our model may therefore be caused by these peptidergic neurons. Orexin (goldfish, zebrafish; [100, 86]) and galanin (*gymnotiform*, [205]) neurons are located in the lateral hypothalamus and project to a region of the ventral telencephalon homologous to the mammalian basal ganglia [72]. The basal ganglia appears to be the core telencephalic region essential for initiating specific motor output [67] and, in both mammals [67] and *gymnotiform* [60, 192, 43], the basal ganglia receives input from dorsal telencephalic recurrent excitatory neural networks. We therefore hypothesize that i) the dorsal telencephalon contains the recurrent networks responsible for the stochastic switching and these networks drive the ventral telencephalon; ii) movement is initiated by the ventral telencephalon; iii) the activity of the ventral telencephalon is modulated, on a slow time scale, by input from the orexin and galanin neurons of the lateral hypothalamus. This may not be an entirely open loop system because the lateral hypothalamus of *gymnotiform* itself receives a strong input from dorsal telencephalon [60]. Further work into the details of the entire network will, however, be required to elucidate how such descending input might regulate behavioural state transitions.

2.4.2 Limitations of the Model

As our goal is to capture the core phenomenology of the Jun et al. [95] data with a minimal set of assumptions, some features of the data inevitably fall outside the scope of our modelling framework. Notably, we do not attempt to reproduce the increased EODR fluctuations in the active state (Figure 4), nor the variable intensity of these fluctuations, as quantified by f_4 (Figure 5A). Although implementing a mapping between the latent variable and the width of the active state well could resolve this discrepancy, we deem the available data as insufficient to rigorously constrain such an ad hoc addition to the model. A biophysically-motivated description of the neural dynamics would be required to properly address inquiries into these finer details of the data. To our knowledge, no such modelling efforts have been made in the context of intermittent locomotion. However, given the apparent two-state and stochastic nature of the EODR, we can draw parallels with the phenomena of cortical Up- and Down-states, which also show higher levels of fluctuations in the Up-state [202].

Models of Up-Down state transitions typically comprise an excitatory and a regulatory component (e.g., inhibition, short term depression, or adaptation), also known as activator/repressor dynamics [76]. The interplay between both components establishes bistability in the network activity and allows stochastic switching between stable states once noise is added to the system [83, 130]. Although these Up-Down state transitions occur at around 0.5 - 2 Hz, much shorter durations than those of the behavioural state transitions examined in our paper, similar dynamical principles might be able to explain both phenomena. For instance, to explain the enhanced fluctuations in Up-states, Hidalgo et al. (2012) suggest a general mechanism that entails stochastic perturbations of the system around the Up-state, for which the associated fixed point is a stable focus [76]. This results in the amplification of a resonant frequency that is absent from the Down-state, where the fixed point is rather a stable node. The idea to represent the Up-state as a stable focus has also been developed for a purely deterministic system [59]. These studies suggest that adding a repressor variable to our model, in such a way as to have the active state fixed point become a focus (which requires two-dimensional dynamics locally), might be sufficient to capture this aspect of the Jun et al. [95] data.

Adding a second dimension to the model might also help to capture the asymmetry observed in the shape of individual active states: most of them show a slow downward

trend long before the movement offset time, followed by an active-to-inactive transition period that is less abrupt than that of the inactive-to-active transitions (Figure 3, insets, and Figure 6). We also speculate that, in a two-dimensional context, differences in timescales for the activator and repressor variable might cause the path toward the active state to be different from the one leaving it. Although a two-variable system might be needed to achieve this asymmetry, we believe that the one-dimensional analyses carried out in this paper remain valid in the sense that trajectories to and from the active state might be governed by distinct one-dimensional dynamics local to each branch (see Pikovsky and Kurths [155] for an example of how a 1-D description can approximate part of a stochastic 2-D trajectory). This is why we analysed transitions around movement offset independently of those around movement onset.

2.4.3 Spontaneous Movement as Decisions Making

Because movement initiation fundamentally emerges from decision-making processes [174], quantitative models of decision-making might provide an adequate substrate to understand the randomness observed in intermittency. So-called accumulator models, for instance, have been extensively and successfully applied to perceptual decision-making tasks [199]. Yet, their validity with respect to spontaneous movements has only been suggested recently in Schurger et al. [173]. The authors of that study contend that the well-known readiness potential [105, 116] that precedes the initiation of movement reflects ongoing spontaneous fluctuations in neural activity, rather than activity related to motor preparation and planning, as traditionally believed. To support this claim, they show that a 1-D model, the leaky stochastic accumulator, with a threshold applied to its output, provides a satisfying fit to EEG data from spontaneous movement tasks. In formal terms, their model is a biased Ornstein-Uhlenbeck process, which describes the evolution of a Brownian variable in a parabolic potential landscape. Although our model variable evolves in a double well potential landscape, the trajectory near either attractor states does approximate an Ornstein-Uhlenbeck process. With the unstable point of the double well acting as a threshold, the model we present here is therefore in the same spirit as that of Schurger et al. [173]. We thus believe that our contribution supports their conclusions, and that spontaneous movement initiation in electric fish is also governed by ongoing fluctuations, with threshold crossing as the mechanism triggering the initiation of movement.

2.4.4 Long tail RTDs in Intermittent Behaviour

A large variety of spontaneous animal behaviour, including intermittent locomotion, is known to be scale-free with power-laws describing the duration distributions of certain behaviours [71, 127, 11]. It has been shown in Proekt et al. [158] that this is a generic trait of a broad class of non-equilibrium Markovian systems. According to this study, the essential requirement is the existence of a macroscopic timescale that imposes an upper bound on the duration of behavioural states and beyond which scale invariance breaks down, which they suggest might arise from species-specific metabolic processes. They noted that the ubiquity of scale invariance in spontaneous behaviour suggests the existence of "an elementary un-differentiated process in the nervous system that governs activation of all behaviors [sic]". We believe the model presented herein provides such a general mechanism: by continuously modulating the switching rates between attractor states, the non-stationary latent variable, interacting with neural noise, is read out as a behavioural driver that evolves over a wide range of timescales, giving rise to the long tail nature of the RTDs. We speculate that the slower of these timescales, that of the latent variable itself, fulfills the role of the macroscopic time scale suggested in Proekt et al. [158].

In the case of electric fish, the RTDs are stretched exponentials, but our modelling framework could also generate power-law RTDs. Indeed, it has been shown that a stochastic bistable system with a non-stationary energy barrier, similar to the model we present here, can yield power-law RTDs in certain regimes [193]. This emergence of power-law and stretched exponential distributions in non-stationary bistable systems can be understood in light of the fact that i) continuously modulating the energy barrier introduces a distribution of switching rates over the course of the experiments, ii) individually, these switching rates translate into exponentially distributed residence times [55], and iii) it has been shown that sums of exponentials can approximate long tail distributions [49, 117, 92]. In addition, an approach similar to ours has been effective in explaining the appearance of power-law and long-tailed RTDs in transitions between cortical Up- and Down-states [130].

2.4.5 Future Work

As *in vivo* recordings of telencephalic activity during spontaneous behaviour become available, it should be possible to constrain a biophysically motivated neural network model to describe the action selection circuitry subserving movement initiation and termination. This could help elucidate the precise mechanism by which bistable attractor dynamics is established in the network, thereby exposing a correspondance between biophysical parameters and the potential function introduced in this paper.

Genetic manipulation of orexin and galanin neurons of the zebrafish hypothalamus may also directly test whether these peptides provide the slow modulatory variable that controls state switching. This could yield a biophysical model of the slow latent variable, which would be used as input to the bistable network. This could help understand the neurophysiological basis of how neuromodulation affects the potential function.

Furthermore, the analyses presented here could be refined by considering a richer behavioural space than the binary, active-inactive classification that we used here. With continuous video tracking, other behavioural states can be used to quantify fish behaviour, such as backward swimming and turning. Based on the transition probability matrix obtained from this description, it might be possible to propose a stochastic multistable attractor model to describe more fully the observed spontaneous behaviour.

Finally, although we have exclusively focused on spontaneous behaviour, it remains to be understood how intrinsic behavioural drivers interact with sensory inputs, i.e., what sort of mechanism implements the influence of sensory perturbations? Is it mediated only through neuromodulation, as suggested by the induced Up-states of the telencephalon [43], or is there a more direct pathway interacting with the inferred attractor dynamics? In other words, one might want to eventually include closed loop features - going beyond the open-loop perspective that underlies the work presented here - to gain a better understanding of how behavioural state transitions occur in more complex settings. Our work also provides an adequate framework to understand sensory responsiveness: depending on the configuration of the potential function and the height of the energy barrier, fish might be more or less inclined to undergo a behavioural state transition following sensory input.

Data-driven jump-diffusion modelling with application to membrane voltage fluctuations in pyramidal neurons

Abstract

The emergent activity of biological systems can often be represented as low-dimensional, Langevin-type stochastic differential equations. In certain systems, however, large and abrupt events occur and violate the assumptions of this approach. We address this situation here by providing a novel method that reconstructs a jump-diffusion stochastic process based solely on the statistics of the original data. Our method assumes that these data are stationary, that diffusive noise is additive, and that jumps are Poisson. We use threshold-crossing of the increments to detect jumps in the time series. This is followed by an iterative scheme that compensates for the presence of diffusive fluctuations that are falsely detected as jumps. Our approach is based on probabilistic calculations associated with these fluctuations, and on the use of the Fokker-Planck and the differential Chapman-Kolmogorov equations. After some val-

idation cases, we apply this method to recordings of membrane noise in pyramidal neurons of the electrosensory line lobe of weakly electric fish. These recordings display large, jump-like depolarization events that occur at random times, the biophysics of which is unknown. We find that some pyramidal cells increase their jump rate and noise intensity as the membrane potential approaches spike threshold, while their drift function and jump amplitude distribution remain unchanged. As our method is fully data-driven, it provides a valuable means to further investigate the functional role of these jump-like events without relying on unconstrained biophysical models.

3.1 Introduction

Complex systems are ubiquitous in many areas of science, including biology, neuroscience, climatology, engineering, as well as in finance and social sciences [69, 84]. The common feature uniting such vastly different systems is the nonlinear interaction between their numerous microscopic constituents. This collective activity leads to the emergence of macroscopic order that cannot be reduced to microscopic properties [145]. It is often these macroscopic variables that can be measured experimentally, allowing the emergent dynamics of the system to be captured by low-dimensional data sets. The search for a macroscopic-level representation of the system thus relies on extracting dynamical information from observed time series [29].

For some systems, the high-dimensional, microscopic degrees of freedom can be well-approximated by simple stochastic fluctuations. These fluctuations often participate as dynamical noise in the macroscopic evolution of the system. The low-dimensional representation of the system's dynamics can then be expressed as a stochastic dynamical system [54]. More specifically, the observed data are usually assumed to satisfy a Langevin-type stochastic differential equation (SDE). A common approach is to obtain the drift and diffusion functions of this equation by estimating the first and second Kramers-Moyal coefficient [177, 110]. The resulting model is completely data-driven, and captures the core phenomenology of the original data without relying on knowledge or assumptions about the microscopic constituents of the observed system. This approach has been successfully applied in a variety of context, ranging from neuronal dynamics [196, 159, 37, 79, 78], heart rate variability [153, 154], turbulence [160, 195], calibration of optical tweezers [140], and others (see Ref. [54] for a review).

In all mentioned cases, however, the noise is assumed to be purely diffusive, i.e., random fluctuations with continuous sample paths. This description is incomplete if, in addition to diffusive fluctuations, large and abrupt events appear at random times throughout the time series. In this case, jump-diffusion stochastic processes provide a more appropriate framework to model these data. Jump-diffusion processes have been used in neuroscience as a model for the spatial [142, 63, 64, 178], and temporal [179] organization of synaptic bombardment, in physics as a model for noise-driven transport in ratchet potentials [120, 121, 34], as well as in finance [106, 190], and soil moisture dynamics [35]. The Langevin approach is likely to fail if the observed system exhibits jump-diffusion characteristics, such as skewed distributions and sudden large jumps. In such cases, and especially when the microscopic dynamics are unknown, extracting a phenomenological model from the experimental data would provide a valuable tool to probe the dynamics of the observed system, its interaction with other systems, and the interplay between diffusive and jump noise sources in shaping the observed behaviour.

In this paper, we present a novel, data-driven inference method that generates an effective jump-diffusion SDE of experimental time series. By detecting jumps through threshold-crossing, and by calculating the contribution of diffusive fluctuations that are falsely detected as jumps, we iteratively estimate the drift function, noise intensity, jump rate, and jump amplitude distribution. The result of this method is a semi-parametric jump-diffusion SDE that successfully fits the original data. Our method is applicable in cases where these data are stationary, with additive diffusive noise and Poisson jumps. Note that other studies have attempted to infer jump-diffusion dynamics from data, but they rely on assuming a parametric form of the jump amplitude distribution [8, 91, 3], or consider only Lévy processes [113, 51].

We test our method with two validation cases, where realizations of known jump-diffusion processes with different characteristics are used as validation data. In both cases, by using only the simulated time series, we precisely recover the correct parameters and functions used in the original simulations. We also compare the autocorrelation functions (ACF) of the validation data and of effective SDE. We then apply our method to recordings of intrinsic membrane voltage fluctuations in pyramidal neurons of electric fish. These recordings contain sudden and unpredictable jump-like events that occur among more typical diffusive membrane noise. Although the exact biophysical

origin of these fluctuations is unknown, we find that the recordings are well-fitted by jump-diffusion processes. We evaluate goodness of fit quantitatively by comparing the observed and the estimated probability density functions (PDF), and power spectral densities (PSD). Interestingly, we find that some pyramidal cells increase their jump rate and noise intensity as the membrane potential approaches spike threshold, while their drift function and jump amplitude distribution remain unchanged.

In Sect. 3.2, we present the various step involved in the inference procedure, including the detection scheme, the calculation relating to false positive, and the iterative approach. In Sect. 3.3, we validate this procedure against one pure diffusion and two jump-diffusion test cases, and then apply it to neurophysiological recordings of membrane noise. This is followed by a discussion on the possible generalizations of our method, and on future work with the experimental data (Sect. 3.4).

3.2 Methods

3.2.1 Definitions and overview

Let $\{X(t)\}$ represent data in the form of a stationary time series. In what follows, $\{X(t)\}$ is obtained either from experimental observations or from numerical simulation. In either case, the situation of interest here is when $\{X(t)\}$ exhibits both diffusive fluctuations and abrupt events (henceforth called jumps). Our goal is then to generate an effective representation of $\{X(t)\}$ in the form of a jump-diffusion SDE:

$$dY(t) = F(Y(t))dt + \sqrt{2D} dW(t) + dJ(t), \quad (47)$$

where F is the drift function, D is the noise intensity, and $W(t)$ is a Wiener process (i.e., Brownian motion). Here $J(t)$ is a compound Poisson process representing the jumps:

$$J(t) = \sum_{i=0}^{N_\lambda(t)} B_i, \quad (48)$$

where $N_\lambda(t)$ is a Poisson point process with rate λ , and the B_i 's are the independent and identically distributed jump amplitudes drawn from a distribution Q_B . For a small enough sampling interval, or time step, Δt , jumps will occur with probability $\Gamma_B \equiv \lambda\Delta t$.

Although the method developed herein is applicable to a wide array of experimental data types, we do, however, impose certain conditions on the underlying dynamical process. Notably, we limit our analysis to systems where the dynamics are stationary (in the strict sense), the diffusion noise is additive, the jumps have positive amplitudes ($B_i > 0$), and where the Poisson rate λ is constant in time, and small enough so that $\Gamma_B \ll 1$. Furthermore, we assume here that $F(x)$ is continuous and that a single stable fixed point arises from the deterministic part of the dynamics, but our method could be generalized to multistable systems. Except for this restriction, we assume no particular shape for the drift function, as long as it generates a stationary process. Finally, we assume that, on average, jump amplitudes are similar, or greater than the typical magnitude of diffusive increments: $E[B_i] > (2D\Delta t)^{\frac{1}{2}}$.

In terms of the time series itself, we assume that the data were sampled at a high-frequency, such that Δt can be assumed small with respect to the total duration of the time series. Note that the value of Δt is set from the experiments that produced the data, and so is not a variable we can control. For experimental data, however, it is possible that the Markov property breaks down on the timescale of individual observations [54], but we assume that, in that case, the time series can be downsampled to a timescale where the Markov property then holds.

Furthermore, in the cases considered below, jumps appear unambiguously in the data, and so a jump-diffusion approach is warranted. In situations where this is not the case, the presence of jumps can be assessed from the fourth-order Kramers-Moyal coefficient, which is non-vanishing for processes with discontinuities [45, 3]. Lastly, note that we make no assumptions regarding the structure of the ACF of $\{X(t)\}$; we only assume stationarity in these data. In fact, we compare the ACFs (and PSDs) of time series generated by the effective SDE with that of the original data as a means of validating the inference procedure.

Below we develop a data-driven inference procedure that successfully generates the estimates \hat{D} , $\hat{\lambda}$, \hat{F} , and \hat{Q}_B , where the unknown functions F and Q_B are estimated non-parametrically. This inference procedure results in a stochastic process, $Y(t)$, that becomes an adequate representation of the original data. In our calculations, we thus associate the first-order stationary PDF of $Y(t)$ to the empirical PDF of $\{X(t)\}$ (obtained

by kernel density estimation), $P_X(x)$.

Our approach is predicated on the detection of jumps in the data via the application of a threshold, $\theta^* > 0$, on the increments $\Delta X(t) \equiv X(t) - X(t - \Delta t)$. This procedure creates a pool of detected jumps with various amplitudes. Let Q_C be the empirical PDF estimated from these measured jump amplitudes. Also, let n be the total number of increments in the time series, and m be the number of increments whose value is greater than θ^* . We define the (overall) jump detection probability as $\Gamma_C \equiv \text{Prob}\{\text{detecting an increment larger than } \theta^* \text{ across an interval } \Delta t\}$, which we estimate from the data as $\Gamma_C = \frac{m}{n}$.

An inherent challenge with this threshold-crossing approach is that, in addition to the true jumps generated by the compound Poisson process, we also unavoidably detect large diffusive fluctuations that are falsely identified as jumps, henceforth called false positives (FP; Fig. 17, top). A direct estimation of the true jump rate, λ , and of the true jump amplitude distribution, Q_B , is thus impossible because the detected jump pool consists of a mixture of true jumps and false positives. Our main contribution, and the central component of our inference procedure, is the calculation of FP-related statistics, namely, the FP detection probability:

$$\Gamma_A \equiv \text{Prob}\{\text{detecting a diffusive increment larger than } \theta^* \text{ across a given time step } \Delta t\},$$

and the distribution from which FP amplitudes are drawn, Q_A . Once Γ_A and Q_A are calculated, we then extract λ (or, equivalently, Γ_B) and Q_B from Γ_C and Q_C . Note that the subscripts “A”, “B”, and “C” will hereafter refer to FPs, true jumps, and both combined, respectively. More precisely, we measure quantities in “C” from the detected jump pool, we calculate FP statistics in “A”, and we seek the true jump statistics in “B”.

3.2.2 Choice of threshold

To ensure that a minimal number of true jumps are missed during the jump detection procedure (i.e., false negatives), the threshold, θ^* , should ideally be set as low as possible. If set too low, however, the number of FPs becomes so large that the statistics of true jumps, i.e., λ and Q_B , cannot be extracted from the statistical fluctuations of Q_C and Γ_C . We thus aim for an intermediate value of θ^* that captures most true jumps while allowing a manageable number of FPs. This is done by relying on our assumption of positive jump amplitudes and, more precisely, by exploiting the asymmetry between

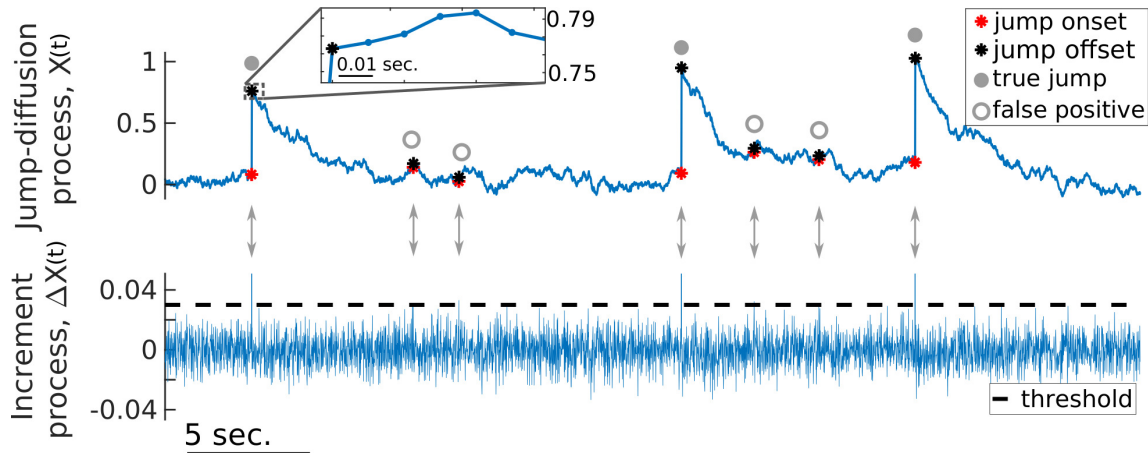


Figure 17: *Jumps are detected by applying a threshold on the increments, but this also creates false positives. **Top:** An example of a simulated jump-diffusion process where both true jumps (grey dots) and false positives (grey rings) are detected. The inset shows how the jump offset can be registered even if subsequent increments are positive. **Bottom:** Increment time series of the simulated process in the top panel. A jump is detected every time an increment exceeds the threshold. In this exaggerated case, true jumps are well above-threshold while false positives barely exceed it. This clear separation is not generally the case.*

positive and negative increment statistics. Note that the threshold depends implicitly on the time step of the original time series: a smaller Δt means that diffusive fluctuations are smaller, and thus a smaller value of θ^* can be used.

Let $\{\Delta X_+\}$ and $\{\Delta X_-\}$ be the sets of positive and negative increments of $\{X(t)\}$, respectively (Fig. 18A). In addition, let $M_+(\theta) = \{\Delta X_+ : \Delta X_+ > \theta\}$ and $M_-(\theta) = \{-\Delta X_- : -\Delta X_- > \theta\}$ be the reduced sets truncated by θ , where θ spans the common range of $\{\Delta X_+\}$ and $\{-\Delta X_-\}$. We use the difference between the sample mean of these sets, $\overline{M_+} - \overline{M_-}$, as a function of θ , to quantify the relative importance of true jumps with respect to diffusive fluctuations, over different increment sizes. The value of θ for which $\overline{M_+} - \overline{M_-}$ is a maximum corresponds to the greatest separability between positive and negative increments. We find, however, that using the inflection point, i.e., where the second derivative becomes 0, located to the left of this maximum (Fig. 18B, asterisks) is a better choice of threshold. This slightly lower value retains a greater range of true jumps, which is desirable, while avoiding the inclusion of an overwhelmingly large number of FP. Choosing this inflection point rather than the maximum impacts primarily the estimation of λ , since it relies on the proper detection of true jumps in the time series. A slightly higher value of θ^* will introduce a bias in $\hat{\lambda}$. For instance, in the first jump-diffusion validation case that we consider in Sect. 3.3.2, choosing the maximum

as the threshold roughly yields a 1% increase in the error on $\hat{\lambda}$ compared to when the inflection point is used. Our approach for setting θ^* is thus motivated by the fact that it is advantageous to choose a value as small as possible for θ^* , and the inflection point in the curve of $\overline{M}_+ - \overline{M}_-$ provides a way reliable way to achieve this.

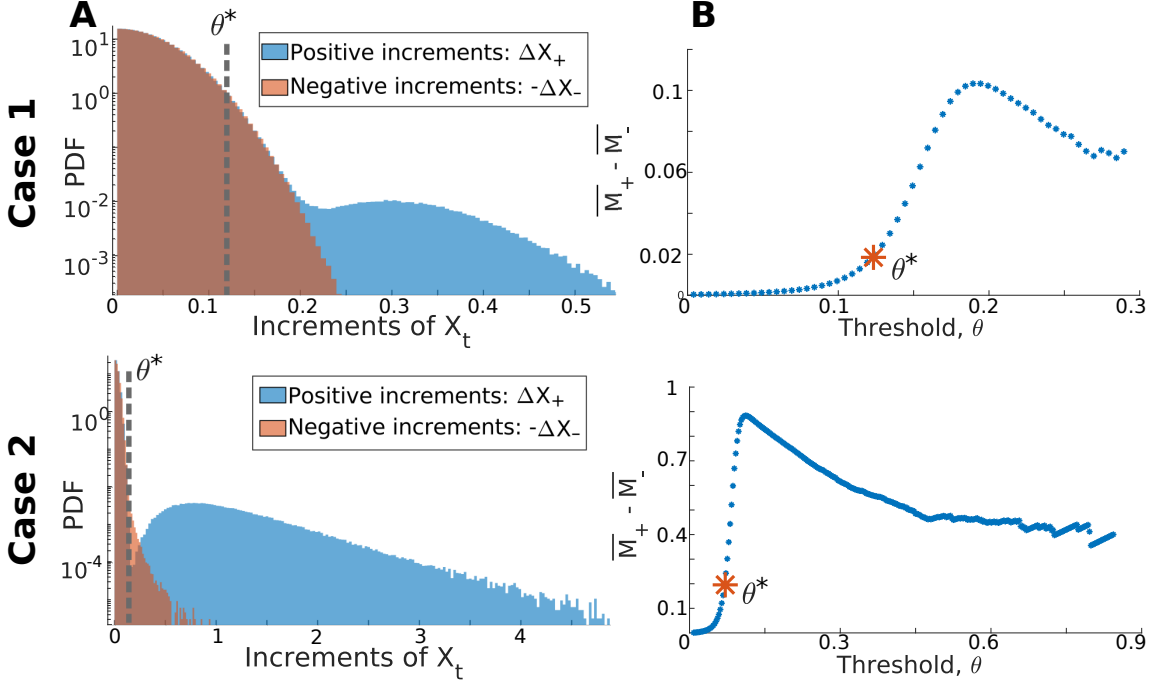


Figure 18: *The choice of θ^* is based on comparing the statistics of positive and negative increments.* We show here how this strategy is applied to the two jump-diffusion validation cases presented in Sect. 3.3.2. **A)** The presence of true jumps allows the PDF of positive increments (blue histogram) to be differentiated from that of negative increments (red histogram) above a certain threshold (dashed line). **B)** This threshold is chosen as the inflection point (red asterisks) of the difference between the sample means of the (truncated) positive and negative increments.

3.2.3 Jump detection

Here we describe how the threshold is applied to the increments in order to generate the detected jump pool. We apply a detection scheme tailored specifically to handle two aspects that we observe in experimental data with jumps, and is inspired by the method used in [97]. Firstly, if the data are resolved on a fine enough time scale, jumps may last longer than a single sampling interval. Secondly, jumps need not be followed immediately by negative increments. In data, and in some simulations as well, the diffusive increments following a jump can still be positive, but we seek a method for identifying when the

jump actually ends (Fig. 17, inset). These two considerations shape the method used to calculate the FP amplitude distribution in Sect. 3.2.4.

When a given increment is larger than the threshold, a jump onset time, T_{on} , is registered, and if the next increment is below threshold, the associated offset time, T_{off} is registered (even if this increment is positive). This defines a jump of duration $T_{\text{off}} - T_{\text{on}} = \Delta t$, where Δt is the sampling interval of the data. We henceforth refer to this type of jump as a singlet, as it spans the duration of a single time step. In contrast, if two or more successive increments are above-threshold, then the jump is of duration $2\Delta t$ or more, which we refer to as a doublet, triplet, and so forth. In other words, the jump offset time is only registered at the end of the sequence of above-threshold increments. With the onset and offset times identified, a jump amplitude is defined as the difference between $X(T_{\text{off}})$ and $X(T_{\text{on}})$.

3.2.4 FP and true jump statistics

Here we present the calculations of the FP detection probability, Γ_A , and of the FP amplitude distribution, Q_A . We then show how these quantities are used to extract the true jump rate, λ , and the true jump amplitude distribution, Q_B from the detected jump probability, Γ_C , and the detected jump amplitude distribution, Q_C . For the calculations in this subsection, we assume that the drift function, F , and noise intensity, D , are known. In the next subsection, we show how these calculations can be incorporated into an iterative scheme that allows the simultaneous estimation of all unknowns, including F and D . Furthermore, note that the FP-related calculations involve only the diffusive part of Eq. (47) since, by definition, FPs occur during the purely diffusive segments between true jumps. We perform validation tests of the calculations of Γ_A and Q_A in Sect. 3.3.1.

FP detection probability

As mentioned in Sect. 3.2.1, sampling a jump-diffusion process at finite intervals and applying a threshold on the observed increments leads to the detection of FPs, i.e., diffusive (rather than true jump) increments larger than the threshold. Importantly,

these FPs occur with a probability that depends on the value of the process at the start of the interval. Let this detection probability be defined as

$$\begin{aligned}\alpha(x) &\equiv \text{Prob}\{\text{detecting an FP in the interval } [t, t + \Delta t], \text{ given that } X(t) = x\} \\ &= \text{Prob}(X(t + \Delta t) - X(t) > \theta^* | X(t) = x),\end{aligned}\tag{49}$$

As α does not depend explicitly on time, this definition relies on our assumption that $\{X(t)\}$ is stationary. The x -dependence arises from the drift function, $F(x)$. Indeed, if the drift function is positive (respectively, negative) at a given time, it biases diffusive fluctuations toward (away from) the threshold. This translates into an FP detection probability that assumes higher values when $F(x) > 0$ than when $F(x) < 0$. We now turn to the explicit calculation of $\alpha(x)$.

Let an increment be defined as $\Delta X(t+\Delta t) = X(t+\Delta t) - X(t)$ and let $\Xi_{\Delta X|X}(\xi|X(t) = x)$ denote its PDF conditioned on the value of the process at the start of the interval, and where ξ assumes the possible values of the increments. Note that, because the time step remains constant, it is always implied that the increments are defined across an interval Δt . Given that we assume Δt sufficiently small, we approximate $\Xi_{\Delta X|X}$ as the short-time propagator of the Fokker-Planck equation [61, 165] associated with the diffusive part of jump-diffusion process (recall that what concerns us here are the purely diffusive segments between the true jumps of $\{X(t)\}$). We thus have $\Delta X(t + \Delta t) \sim \mathcal{N}(F(X(t))\Delta t, 2D\Delta t)$, and

$$\Xi_{\Delta X|X}(\xi|X(t) = x) \approx \frac{1}{\sqrt{4\pi D\Delta t}} \exp\left(-\frac{(\xi - F(x)\Delta t)^2}{4D\Delta t}\right),\tag{50}$$

that is, a Gaussian distribution with mean $F(x)\Delta t$ and variance $2D\Delta t$.

For the test cases presented in Sect. 3.3 (with $\Delta t = 10^{-4}$ s), we have validated this approximation by comparing it with numerical solutions of the associated Fokker-Planck equation solved at a finer temporal resolution ($\Delta t/1000$) across the time step Δt . The numerical solutions were indeed well-fitted with the approximation in Eq. (50) (not shown). Numerical integration was performed with a custom partial differential equation solver that implements a finite volume discretization along with the fully implicit Euler scheme. The advective term was treated with the upwind scheme and a linear interpolation profile for the spatial derivative was applied to the diffusive term. The resulting algebraic equation was solved with the tridiagonal matrix algorithm [150].

Once the conditional PDF of the increments is evaluated with Eq. (50), we calculate the x -dependent FP detection probability as:

$$\alpha(x) = \int_{\theta^*}^{\infty} \Xi_{\Delta X|X}(\xi|X(t) = x)d\xi, \quad (51)$$

that is, the probability of observing an increment larger than θ^* starting at x . Finally, the detection probability that we seek is the average:

$$\Gamma_A = \mathbb{E}[\alpha(X(t))] \quad (52)$$

$$= \int_{-\infty}^{\infty} \alpha(x)P_X(x)dx. \quad (53)$$

We validate these calculations in Sect. 3.3.1.

FP amplitude distribution

We now proceed with the calculation of Q_A , i.e., the distribution from which FP amplitudes are drawn. First, recall that our detection scheme allows for jumps of different durations (Sect. 3.2.3). As such, the detection of an FP implies either a succession of above-threshold increments (e.g., Fig. 19A), or, at least, a single above-threshold increment. Let T_{on} denote the FP onset time, that is, the time at the start of the first above-threshold increment, and let $X_0 \equiv X(T_{\text{on}})$ be referred to as the starting value of the FP. Moreover, let τ denote the FP duration, an integer multiple of Δt , such that $\tau = \Delta t$ corresponds to an FP singlet, $\tau = 2\Delta t$ to an FP doublet, and so forth.

In order to calculate Q_A , let us first identify the factors that influence FP amplitudes. Firstly, it must be noted that the FP amplitudes will exhibit a similar x -dependence as that discussed in the preceding section. Indeed, an increment starting at $X_0 \equiv X(T_{\text{on}}) = x$ will tend to be larger when $F(x) > 0$ than when $F(x) < 0$. Secondly, the amplitude of an FP will also depend on its duration, τ . For instance, the three increments of a triplet FP will summate and tend to have a larger amplitude than that of a singlet. The FP amplitudes, A , will thus covary with X_0 and with τ , but note that τ also depends on X_0 . Indeed, longer FPs will tend to occur where the drift function is more positive, and vice versa. To account for these dependencies, let us define the trivariate random variable $\{A, \tau, X_0\}$, distributed according to its joint PDF $P_{A,\tau,X_0}(a, i\Delta t, x)$, where we explicitly

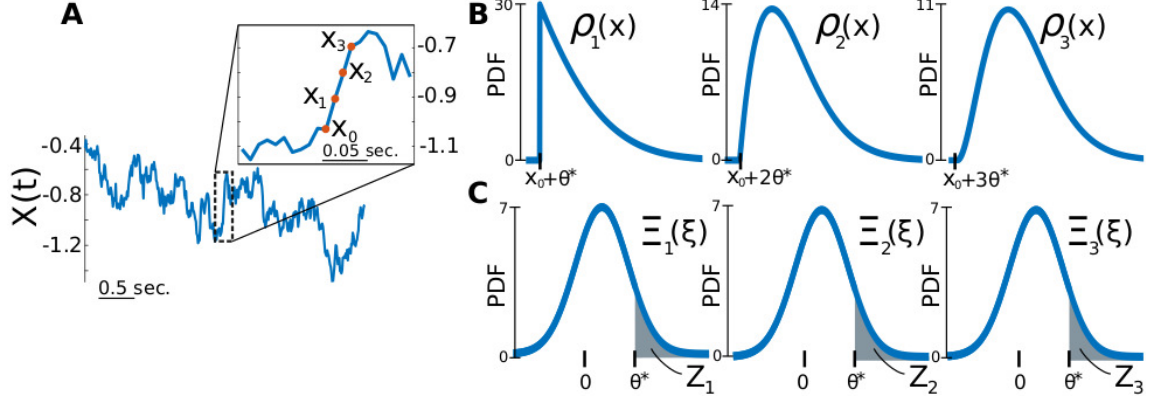


Figure 19: The estimate of Q_A is obtained through a probabilistic analysis of FP detection. **A)** Example of a diffusive fluctuation registered as an FP triplet. **B)** The situation in A is addressed by calculating ρ_i , the PDF of X_i conditioned on X_0 . **C)** From ρ_{i-1} , we then calculate Ξ_i , the PDF of ΔX_i conditioned on X_0 , which is used to evaluate Z_i (shaded area), the probability that the i^{th} increment is above threshold, given X_0 . Note that, although they look similar, the Ξ_i 's are slightly different from each other.

write τ as an integer multiple of Δt . What we seek then is the marginal:

$$Q_A(a) = \int_{-\infty}^{\infty} \sum_{i=1}^{\infty} P_{A,\tau,X_0}(a, i\Delta t, x) dx, \quad (54)$$

where the sum extends over all possible FP durations. From the definition of conditional PDFs, we can expand the joint PDF as

$$\begin{aligned} P_{A,\tau,X_0}(a, i\Delta t, x) &= P_{A|\tau,X_0}(a|i\Delta t, x) P_{\tau,X_0}(i\Delta t, x) \\ &= P_{A|\tau,X_0}(a|i\Delta t, x) P_{\tau|X_0}(i\Delta t|x) P_{X_0}(x), \end{aligned} \quad (55)$$

where $P_{\tau|X_0}(i\Delta t|x) = \text{Prob}(\text{detecting an FP of duration } i\Delta t, \text{ given the starting value } x)$ is the conditional probability mass function of the FP duration τ . We can thus write Eq. (54) as⁶

$$Q_A(a) = \int_{-\infty}^{\infty} \left(\sum_{i=1}^{\infty} P_{A|\tau,X_0}(a|i\Delta t, x) P_{\tau|X_0}(i\Delta t|x) \right) P_{X_0}(x) dx. \quad (56)$$

The sum in the large parentheses is a function of x , and Eq. (56) is merely calculating its average with respect to the starting value X_0 . This sum can further be interpreted as a so-called mixture distribution: consider a collection of random variables, one of which is chosen according to a certain probability (its *mixture weight*), and is then realized

⁶Alternatively, one could obtain Eq. (56) directly from Eq. (54) by invoking the law of total probability [149] for the set of mutually exclusive and exhaustive events $\{\tau = i\Delta t, i \in \mathbb{N}^*\}$.

according to its own PDF (its *mixture component*). The outcome of this experiment is itself a random variable whose PDF is called a *mixture distribution*, and is expressed as a sum over the PDFs of the random variables in the collection, weighted by their respective probabilities. In our case, for a fixed value of X_0 , an FP duration is drawn according to a countable set of mixture weights, $P_{\tau|X_0}$, and an FP amplitude is then realized according to the associated mixture component, $P_{A|\tau,X_0}$. In practice, the sum will be truncated after the first few terms because the subsequent mixture weights become negligible.

From Eq. (56), we see that in order to arrive at the desired Q_A , the functions $P_{A|\tau,X_0}$, $P_{\tau|X_0}$, and P_{X_0} must first be calculated. Let us first consider the latter. Because X_0 represents, by definition, the value of $X(t)$ at the start of an above-threshold increment, we can express its PDF in terms of the joint PDF of $\Delta X(t + \Delta t)$ and $X(t)$:

$$\begin{aligned}
P_{X_0}(x) &= K \int_{\theta^*}^{\infty} P_{\Delta X, X}(\xi, x) d\xi \\
&= K \int_{\theta^*}^{\infty} \Xi_{\Delta X|X}(\xi|X(t) = x) P_X(x) d\xi \\
&= K P_X(x) \int_{\theta^*}^{\infty} \Xi_{\Delta X|X}(\xi|X(t) = x) d\xi \\
&= K P_X(x) \alpha(x),
\end{aligned} \tag{57}$$

where K is a normalization constant, and where we integrate from θ^* to enforce that X_0 is associated with the onset of an above-threshold increment. Equation (57) tells us that the starting value of an FP, X_0 , is distributed according to a modified version of P_X that accounts for the fact that FPs are not equally likely to occur everywhere along the x -axis. Let us now consider the calculation of $P_{A|\tau,X_0}$.

In what follows, we simplify the notation by labelling time with the index i , such that $i = 0$ represents the time T_{on} , $i = 1$ the time $T_{\text{on}} + \Delta t$, $i = 2$ the time $T_{\text{on}} + 2\Delta t$, and so forth. With this notation, $X_i \equiv X(T_{\text{on}} + i\Delta t)$ denotes the i^{th} point following the FP onset time, and $\Delta X_i \equiv X_i - X_{(i-1)}$ the i^{th} increment. Furthermore, our focus here is on FPs of duration, say, $i\Delta t$, which corresponds to a sequence of i successive above-threshold increments. As such, the forthcoming calculations involve PDFs that are implicitly conditioned on the event $\{\Delta X_n > \theta^*, \forall n \leq i\}$.

For an FP of duration $\tau = i\Delta t$ starting at X_0 , we define its amplitude as $A = X_i - X_0$, and we seek the conditional PDF $P_{A|\tau,X_0}$. For this purpose, let $\rho_i(x) \equiv P_{X_i|X_0}(x|x_0)$

denote the PDF of X_i , $i > 1$, conditioned on X_0 , and let $\rho_0(x) = \delta(x - x_0)$ denote the PDF of X_0 , which we treat here as a fixed initial condition⁷. Since A is expressed as a difference, we can directly write

$$\begin{aligned} P_{A|\tau,X}(a|i\Delta t, x_0) &= \int_{-\infty}^{\infty} P_{X_i, X_0}(x + a, x) dx \\ P_{A|\tau,X}(a|i\Delta t, x_0) &= \int_{-\infty}^{\infty} \rho_i(x + a) \delta(x - x_0) dx \\ P_{A|\tau,X}(a|i\Delta t, x_0) &= \rho_i(a + x_0). \end{aligned} \quad (58)$$

The ρ_i 's, for $i > 1$, are evaluated sequentially based on the fact that $X_i = \Delta X_i + X_{i-1}$. The PDF of this sum, conditioned on X_0 , gives

$$\begin{aligned} \rho_i(x) &= \int_{-\infty}^{\infty} P_{\Delta X_i, X_{i-1}|X_0}(\xi, x - \xi|x_0) d\xi \\ &= \int_{-\infty}^{\infty} P_{\Delta X_i|X_{i-1}, X_0}(\xi|x - \xi, x_0) P_{X_{i-1}|X_0}(x - \xi|x_0) d\xi \\ &= \int_{-\infty}^{\infty} P_{\Delta X_i|X_{i-1}}(\xi|x - \xi) \rho_{i-1}(x - \xi) d\xi. \end{aligned} \quad (59)$$

To enforce the condition of above-threshold increments, $\{\Delta X_n > \theta^*, \forall n \leq i\}$, we evaluate $P_{\Delta X_i|X_{i-1}}$ based on Eq. (50), but we truncate the distribution below $\xi = \theta^*$:

$$P_{\Delta X_i|X_{i-1}}(\xi|x) = K \begin{cases} \Xi_{\Delta X_i|X_{i-1}}(\xi|x) & \text{if } \xi > \theta^* \\ 0 & \text{otherwise} \end{cases} \quad (60)$$

where K is a normalization constant. With Eq. (60) and (59), we finalize the calculation of $P_{A|\tau,X}$ in Eq. (58). In Fig. 19B, we see the representation of the ρ_i for an FP triplet. From ρ_i , we can also calculate the PDF of ΔX_i , conditioned on X_0 (this will be useful in the calculation of $P_{\tau|X_0}$). Let this PDF be defined as $\Xi_i(\xi) \equiv P_{\Delta X_i|X_0}(\xi|x_0)$. We calculate it as a marginal over X_{i-1} :

$$\begin{aligned} \Xi_i(\xi) &= \int_{-\infty}^{\infty} P_{\Delta X_i, X_{i-1}|X_0}(\xi, x|x_0) dx \\ &= \int_{-\infty}^{\infty} \Xi_{\Delta X_i|X_{i-1}}(\xi|x) \rho_{i-1}(x) dx, \end{aligned} \quad (61)$$

where the dependence of $\Xi_{\Delta X_i|X_{i-1}}$ on X_0 disappears because of the Markov property. In Fig. 19C, we see the representation of the Ξ_i 's for an FP triplet.

⁷In the final calculation of Eq. (56), we obviously average over these initial conditions by taking the integral over P_{X_0} .

We now turn to the calculation of $P_{\tau|X_0}$, the probability of observing an FP of duration τ , conditioned on the starting value X_0 . We are interested in the conditional probability of the event $\{\tau = i\Delta t\}$, where i is an integer. This event is equivalent to the intersection of the events $E_1 \equiv \{\Delta X_1 > \theta^*\}$, $E_2 \equiv \{\Delta X_2 > \theta^*\}$, ... and $E'_{i+1} \equiv \{\Delta X_{i+1} \leq \theta^*\}$. In other words, we obtain an FP of duration $i\Delta t$ when the first i increments are above-threshold, but the $(i + 1)^{\text{th}}$ increment is below-threshold.

We can thus expand $P_{\tau|X_0}$ as:

$$\begin{aligned}
P_{\tau|X_0}(i\Delta t|x_0) &= \text{Prob}(\tau = i\Delta t|x_0) \\
&= \text{Prob} \left[E'_{i+1} \cap \left(\bigcap_{n=1}^i E_n \right) \middle| x_0 \right] \\
&= \text{Prob} \left[E'_{i+1} \middle| \left(\bigcap_{n=1}^i E_n \right), x_0 \right] \cdot \text{Prob} \left[E_i \middle| \left(\bigcap_{n=1}^{i-1} E_n \right), x_0 \right] \cdot \\
&\quad \text{Prob} \left[E_{i-1} \middle| \left(\bigcap_{n=1}^{i-2} E_n \right), x_0 \right] \cdot \dots \cdot \text{Prob}(E_2|E_1, x_0) \cdot \text{Prob}(E_1|x_0).
\end{aligned}$$

Let $Z_i(x_0) \equiv \text{Prob} \left[E_i \middle| \left(\bigcap_{n=1}^{i-1} E_n \right), x_0 \right]$, $i > 1$, represent the probability that the i^{th} increment is above-threshold, given that the $i - 1$ previous increments were also above-threshold, and given the starting value x_0 . These Z_i 's can be calculated from the Ξ_i 's of Eq. (61) as (Fig. 19B, shaded area):

$$\begin{aligned}
Z_i(x_0) &= \text{Prob}(\Delta X_i > \theta^* | \Delta X_n > \theta^*, \forall n < i; x_0) \\
&= \int_{\theta^*}^{\infty} \Xi_i(\xi) d\xi.
\end{aligned} \tag{62}$$

We now arrive at the desired probability mass function:

$$P_{\tau|X_0}(i\Delta t|x_0) = \text{Prob}(\tau = i\Delta t|x_0) = (1 - Z_{i+1}(x_0)) \prod_{n=1}^i Z_n(x_0), \tag{63}$$

where we have used the fact that $1 - Z_{i+1}$ is equal to the probability that the $(i + 1)^{\text{th}}$ increment is below threshold, and where we have defined $Z_1(x_0) \equiv \text{Prob}(E_1|x_0) = \alpha(x_0)$, i.e., the probability that the first increment after X_0 be above-threshold. Once Eq. (57), (58) and (63) are evaluated, we apply Eq. (56) to obtain the desired Q_A . Using this approach, we obtain an excellent fit between theory and simulations, as reported in Sect.

3.3.1.

True jump rate

Our estimate of the true jump rate, λ , relies on the knowledge of the overall jump detection probability, Γ_C and on the FP detection probability Γ_A (both defined in Sect. 3.2.1). Recall that we calculate Γ_A from Eq. (52), while we estimate Γ_C directly from the data as m/n , where m is the number of time steps with $\Delta X(t) > \theta^*$ (either from a true jump or an FP) and n is the total number of time steps in the data time series. On the other hand, from the definition of Γ_C we can write:

$$\begin{aligned}
 \Gamma_C &\equiv \text{Prob}\{\text{detecting an increment larger than } \theta^* \text{ across an interval } \Delta t\} \\
 &= \text{Prob}\{(\text{detecting an FP across } \Delta t) \cup (\text{detecting a true jump across } \Delta t)\} \\
 &= \Gamma_A + \Gamma_B - \Gamma_A \Gamma_B \\
 &= \Gamma_A + \lambda \Delta t - \Gamma_A \lambda \Delta t,
 \end{aligned} \tag{64}$$

where $\Gamma_B = \lambda \Delta t$ is the probability of observing a true jump in an interval Δt . This is merely a statement of the addition law of probability, which would read: the probability of detecting a jump in an interval Δt is the sum of the probability of observing a true jump, plus that of observing an FP, minus the probability of observing both at the same time, where we use the fact that FPs and true jumps are independent events. For the test cases described in Sect. 3.3.2, isolating λ in Eq. (64) is accurate up to an error of 0.02%.

True jump amplitude distribution

As in the previous subsection, we obtain an estimate for the true jump amplitude distribution, Q_B , based on the empirical PDF of jump amplitudes measured from the time series, Q_C , and on the calculated FP amplitude distribution, Q_A . Because detected jumps are a mixture of true jumps and FPs, we can write, a priori,

$$Q_C = W_A Q_A + W_B Q_B, \tag{65}$$

where $W_A = \frac{\Gamma_A}{\Gamma_C}$ is the probability that a detected jump is an FP, and $W_B = \frac{\Gamma_B}{\Gamma_C}$ that it is a true jump. The subtlety here is that, contrary to FPs, true jumps are never detected on their own, as they always summate with a diffusive fluctuation. In other words, we never observe the B_i 's directly, but rather the B_i 's plus a diffusive increment. Over a short enough time step, diffusive increments are Gaussian variables and are approximately independent of each other. For the purpose of calculating Q_B , we will thus assume these increments are Gaussian with mean zero and variance $2D\Delta t$: $\Delta X_{\text{diff}} \sim \mathcal{N}(0, 2D\Delta t)$. Properly accounting for the x -dependence of the mean would be more precise, but would require Q_C to be broken down into a family of distributions parameterized by x , which would require a very large number of detected jumps in the data.

Let $\tilde{\Xi}_{\Delta X}$ represent a Gaussian distribution with zero mean and variance $2D\Delta t$. The Q_B in Eq. (65) should thus be replaced by the convolution $\tilde{\Xi}_{\Delta X} * Q_B$. Furthermore, W_A must in fact be reduced by a factor $(1 - \Gamma_B)$ to account for the probability that FPs can occur during the same interval as a true jump. This leads to:

$$Q_C = \frac{\Gamma_A(1 - \Gamma_B)}{\Gamma_C} Q_A + \frac{\Gamma_B}{\Gamma_C} (\tilde{\Xi}_{\Delta t} * Q_B). \quad (66)$$

From this equation, we isolate the convolution term and apply the Basic Deconvolution Algorithm [101] to extract Q_B . Let Y be a measured, convolved signal, where the convolution kernel, h , is known. We seek the intact signal X , such that $Y = (h * X)$. We compute an estimate of X at each iteration as $X_{k+1} = X_k + [Y - (h * X_k)]$, with $X_0 = Y$. The algorithm converges once the correct signal is reached, since the residual between Y and $(h * X_k)$ then becomes zero.

3.2.5 Iterative procedure, noise intensity, and drift function

We now turn to the problem of the simultaneous data-driven estimation of all the unknowns in Eq. (47). To this end, we incorporate the calculations of Sect. 3.2.4 in the iterative scheme depicted in Fig. 20, which consists of three main branches. The first two initial branches, I and II, are independent and are performed only once; this is followed by branch III where the iterations take place. In branch I, the threshold is set (Sect. 3.2.2) and then applied to the time series to yield the detected jump pool, from which Γ_C and Q_C are obtained (Sect. 3.2.3). The noise intensity is estimated in branch II, along with an initial guess for the drift function, \hat{F}_1 , which are both used in the first iteration of branch III to calculate $\hat{\Gamma}_B$ and $\hat{\lambda}$ (Sect. 3.2.4). The last step in branch III uses \hat{D} , $\hat{\Gamma}_B$,

and $\hat{\lambda}$ to estimate the drift function, \hat{F} , which is fed back to the first step of branch III in order to iteratively refine the estimation procedure.

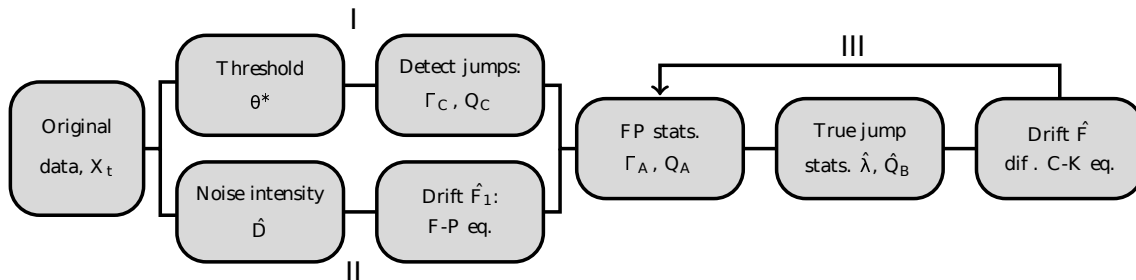


Figure 20: *Overview of the flow of our iterative procedure.* In both validation cases, a satisfying estimate is obtained after about 10 iterations. The threshold and noise intensity are estimated directly from the data in I and II, while the true jump rate and amplitude distribution, as well as the drift function, appear in the iterative phase III. Note that the true jumps statistics are not yet established in II, and this is why we resort to the Fokker-Planck equation as a means to obtain a preliminary guess of the drift function, $\hat{F}_1(x)$. A more refined estimate of $F(x)$ is later obtained at the end of branch III.

Noise intensity

As depicted in Fig. 20, the estimation of D does not rely on the value chosen for θ^* . It does, however, still use the notion of applying a threshold on the increments. Indeed, calculating \hat{D} relies on partitioning $\{X(t)\}$ into mostly jump-free segments, the length and number of which is sensitive to the value of the threshold used for detection: the lower the threshold, the more jumps are detected (some of which are FPs), and the shorter these partitions are, which, as explained below, can skew the estimate of D . A high threshold, on the other hand, leaves a significant number of true jumps in those segments. The goal here is thus to vary the threshold, θ , in order to obtain the optimal estimate of D .

In the limit of an infinitesimally small sampling interval, $\Delta t \rightarrow 0$, the quadratic variation, $[X^{\text{diff}}(t)]$, of a pure diffusion process converges to the so-called integrated variation, which, for additive and time-independent noise, gives [2, 70]:

$$[X^{\text{diff}}(t)] = \int_0^T 2D ds = 2DT, \quad (67)$$

where $T = (n-1)\Delta t$ is the total duration, for the n samples of $\{X(t)\}$. We can, therefore,

estimate D via the sample quadratic variation, also known as realized variance, $RV(t)$ [2, 70]:

$$\hat{D} \approx \frac{1}{2T} RV(t) = \frac{1}{2T} \sum_{k=0}^{n-1} \left(X^{\text{diff}}(t_{k+1}) - X^{\text{diff}}(t_k) \right)^2. \quad (68)$$

For instance, for a test diffusive process with $F(x) = -0.2x$ and $D = 0.15$ and sampled at $\Delta t = 0.01s$ ($N = 10^6$), Eq. (68) estimates D with an error of 0.002%. In contrast, applying Eq. (68) to a realization of a jump-diffusion process returns an overestimated \hat{D} , as expected due to large, non-diffusive, positive increments that populate $\{X(t)\}$. Also note that if the threshold is set low enough to detect the smallest true jumps, in general, it can also remove the largest diffusive increments, which are required to properly apply Eq. (68). Simply removing the detected jumps from the sum in Eq. (68) would, therefore, yield an underestimated \hat{D} .

To circumvent this problem, we consider only the negative increments of $\{X(t)\}$ in the calculation of \hat{D} , as they will remain essentially unaffected by the presence of positive jumps, with the exception of a short transient following the jump offset. Indeed, following each jump, we expect to see a brief period where the process is out of equilibrium. And since the jumps have positive amplitudes, negative increment statistics are biased toward negative values during this transient (e.g., Fig. 17). The calculation of \hat{D} is thus based on applying Eq. (68) to jump-free segments of $\{X(t)\}$, but only including negative increments, and neglecting the initial transient at the start of each segment (the duration of which is determined below). This is repeated for various values of θ . The successful estimation of D based only on negative increments relies on our assumption that Δt is small, for in this case increments are approximately independent and distributed as $\mathcal{N}(0, 2D\Delta t)$. For larger Δt , the increment PDF can become asymmetric, meaning that the statistics of negative increments differ from those of positive ones, which would cause errors in our estimation of D .

Let T_{off} , and T_{on} denote the jump offset and onset times, respectively. Note that the values of these times, and the number of detected jumps all depend on the specific value of the threshold. Then the i^{th} segment is defined by $\{S(t)\}_i = \{X(t) : T_{\text{off}}(i) < t < T_{\text{on}}(i+1)\}$ and is of duration $T_i = T_{\text{on}}(i+1) - T_{\text{off}}(i)$, and let $\{\Delta S(t)\}_i$ be its n_i increments. Out of these n_i increments, we keep only the n_i^- that are negative, and that occur after the transient of approximate duration Φ . We are thus left with the following

subset of increments from each segment:

$$\{\Delta S(t)\}_i^- = \{\{\Delta S(t)\}_i : \{\Delta S(t)\}_i < 0, t > \Phi\}. \quad (69)$$

For each segment, we obtain an estimate \hat{D}_i as:

$$\hat{D}_i = \frac{1}{2T_i^-} \sum_{k=1}^{n_i^-} \left(\{\Delta S(t_k)\}_i^- \right)^2, \quad (70)$$

where $T_i^- = n_i^- \Delta t$ is the effective duration of the combined n_i^- negative increments. We then calculate \hat{D} as an average of the \hat{D}_i 's, weighted by T_i/T .

More precisely, here are the steps taken in order to arrive at \hat{D} :

- Starting from the largest value of $\{\Delta X(t)\}$, lower the threshold until the largest 5% of jumps are detected, which are the ones with the most prominent transient.
- Let $\{S^*(t)\}_i$ be the segments that follow these jumps (Fig. 21A). Average across them for each time step, creating a jump-triggered average trace (Fig. 21A), black line).
- Identify Φ as the approximate moment when the jump-triggered average stabilizes, quantified here as when its derivative is less than 0.05 (results do not depend strongly on this particular value: changing Φ by one order of magnitude on either side of the value used here yields estimates of D that differ by less than 0.2%). This gives us an estimate for the maximum time scale required for post-jump equilibrium.
- With Φ determined, and for each value of the threshold, extract the $\{S(t)\}_i$'s (Fig. 21B) and apply Eq. (70) to their negative increments for which $t > \Phi$ (neglecting segments for which $T_i < \Phi$).

This scheme results in an estimate \hat{D} for each value of the threshold, and the lowest value is chosen as the optimal estimate (Fig. 21C). This is because this approach overestimates D on both ends of the range of threshold values, but for different reasons. For low θ , many more jumps are detected and this makes the $\{S(t)\}_i$'s shorter, which means that the associated \hat{D} is more prone to be biased by the residual of the transient. On the other hand, a very high threshold leaves significant jumps in the $\{S(t)\}_i$'s, which biases the statistics of negative increments. The optimal balance is reached somewhere in between, where the segments are large enough so that the initial transient is negligible, and where the jumps that are inevitably left in the segments do not significantly




Figure 21: Calculating \hat{D} relies on partitioning $\{X(t)\}$ in jump-free segments for different threshold values. We show here how this strategy is applied to the two jump-diffusion test cases presented in Sect. 3.3.2. **A)** A jump-triggered average (black curve) is obtained from the largest 5% jumps in the data and is used to obtain a maximum estimate of the transient time scale Φ (dashed line). **B)** Jump-free segments used to calculate the optimal value of \hat{D} in C. **C)** Different estimates of D are produced for different values of the threshold. We heuristically choose the minimum value as the optimal value. Traces in A and B are from Case 2.

alter the statistics of negative increments. The heuristic choice of an intermediate value, namely one that corresponds to the minimal \hat{D} estimate, gives excellent results in both validation cases (less than 0.1% error, see Table 4).

Drift function

Our estimation of the drift function, $F(x)$, relies on the differential Chapman-Kolmogorov equation [57], which describes the evolution of the PDF of a stochastic process where jumps occur alongside a diffusion term. In the case of positive Poisson jumps, this equation reduces to (see Appendix A):

$$\frac{\partial P_X(x, t)}{\partial t} = -\frac{\partial}{\partial x} [F(x)P_X(x, t)] + D\frac{\partial^2}{\partial x^2} P_X(x, t) - \lambda P_X(x, t) + \lambda \int_0^\infty Q_B(s)P_X(x-s, t)ds. \quad (71)$$

If D , λ , and Q_B are known (or estimated), and if $P_X(x, t)$ is assumed stationary and equal to the empirical PDF of the data at hand, we can invert Eq. (71) to evaluate the drift function.

Note, however, that \hat{F} is required in the first step of branch III of the iterative procedure (Fig. 20), since the FP-related statistics, Γ_A and Q_A , are calculated based on the drift function. A preliminary estimate, \hat{F}_1 , of the drift function is thus required. This particular estimate, which is needed only once throughout the inference procedure, is obtained by letting $\lambda = 0$ in Eq. (71), such that it becomes the Fokker-Planck equation associated with the diffusive part of the stochastic process. The stationary solution of this Fokker-Planck equation can be used to establish a relation between the noise intensity, $P_X(x)$ (the empirical PDF of the data) and the drift function [165, 62]:

$$P_X(x) = \frac{K}{\hat{D}} \exp\left(-\int \frac{\hat{F}_1(x)}{\hat{D}} dx\right), \quad (72)$$

where K is a normalization constant. This first preliminary estimate is necessarily flatter than the true $F(x)$, as the presence of jumps makes $P_X(x)$ wider than it would be if there were no jumps. Successive iterations gradually rectify this by incorporating estimates of λ and Q_B in Eq. (71).

3.3 Results

Here we present three applications of the method developed above. First we validate the calculation of Γ_A and Q_A for the case of a purely diffusive process. Then we apply the full iterative scheme to two simulated jump-diffusion processes with different characteristics. Finally, we apply our inference method to electrophysiological recordings in pyramidal cells of electric fish.

3.3.1 Validation of the FP statistics calculations

To confirm that the calculations of Q_A and Γ_A are accurate, we start with a simple test case where we consider a simulated pure diffusion process, $\{X^{\text{diff}}(t)\}$. As there are no jumps here, the distribution of increments does not possess the necessary asymmetry to properly identify a threshold. For this test case only, we thus opt for a specific value, $\theta^* = 0.1$, that showcases the ability of our method to handle FPs of various durations. The results presented here, however, remain valid for a range of values of θ^* . For the parameters used in this pure diffusion validation case (Table 3), this range extends from 0.025 up to 0.2. The upper limit is set by the fact that, beyond it, too few FPs are de-

tected, which precludes any statistical calculations from being achieved. The lower limit, on the other hand, arises because too many FPs are detected, such that, for instance, they occur every other time step. In such a case, Δt is too large and the estimation of λ becomes imprecise due to the statistical fluctuations in the number of detected FPs.

Applying the threshold in this case leads to a detected jump pool comprised entirely of FPs. The goal now is to compare the measured Q_C and Γ_C , with the calculated, Q_A and Γ_A . If we obtain that $\Gamma_C \approx \Gamma_A$ and that $Q_C \approx Q_A$, then we will effectively have shown that the true jump rate is zero, $\lambda = 0$, and that our method correctly calculates the FP amplitude distribution. In this pure diffusion test case, these calculations rely on the knowledge of the correct noise intensity, D , and the correct drift function, F , but this will not be the case in subsequent sections.

With the particular values of D and F used here to simulate $\{X^{\text{diff}}(t)\}$ (Table 3), we find that FPs are either singlets, doublets, or triplets, which contribute differently to the measured amplitude distribution, Q_C . Indeed, the fact that longer FPs tend to have larger amplitudes, and that FP durations are always multiples of Δt , creates distinct modes in the measured amplitude distribution (Fig. 22A, blue histogram). By taking the sum out of the integral in Eq. (56), we can write $Q_A(a) = \sum_{i=1}^{\infty} Q_i(a)$, where

$$Q_i(a) = \int_{-\infty}^{\infty} P_{A|\tau, X_0}(a|i\Delta t, x) P_{\tau|X_0}(i\Delta t|x) P_{X_0}(x) dx, \quad (73)$$

are the individual distributions associated with FPs of duration $i\Delta t$ (Fig. 22A, yellow curves). These distributions are then summed to obtain Q_A , which is a precise match with Q_C for this purely diffusive example (Fig. 22A, black curve).

Furthermore, by applying Eq. (51) we calculate the x -dependent detection probability, $\alpha(x)$ (Fig. 22B, yellow curve). As expected, this function depends non-trivially on x , and reflects the nonlinearity of the specific drift function used in this example. If multiplicative noise had been used, the noise intensity, $D(x)$, would also have influenced the shape of $\alpha(x)$. To validate this calculation of $\alpha(x)$, we run Monte Carlo simulations of the diffusion process

$$dX^{\text{diff}}(t) = F(X^{\text{diff}}(t))dt + \sqrt{2D} dW(t), \quad (74)$$

over a duration Δt , but with a time step of $\Delta t/1000$, and with various initial conditions

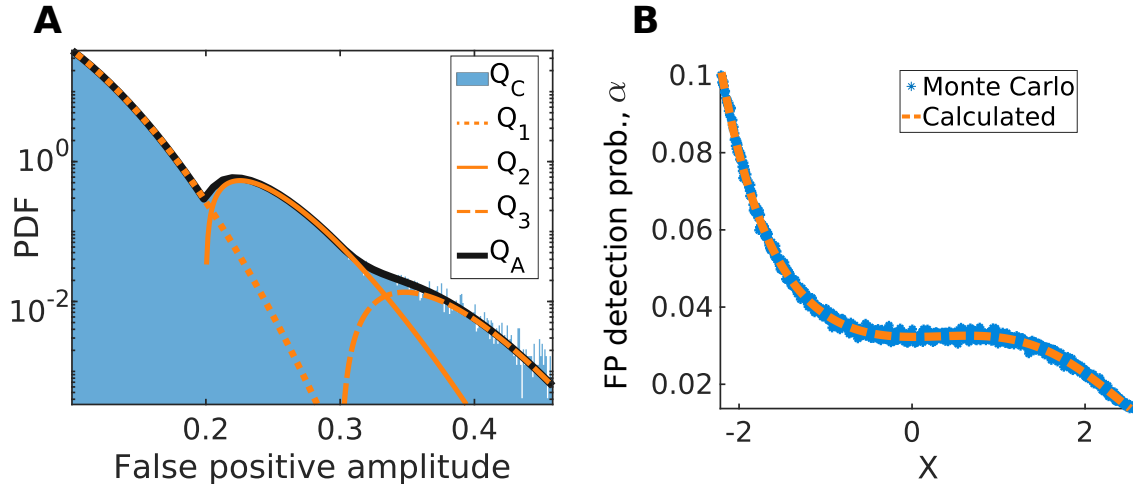


Figure 22: For a pure diffusion process, we correctly calculate the FP amplitude distribution, Q_A , and FP detection probability, Γ_A . **A)** After applying a threshold to $\Delta X^{\text{diff}}(t)$, we obtain a pool of FPs with a range of durations and amplitudes. From the latter we measure Q_C (blue histogram). We correctly calculate this distribution as $\sum_i Q_i$, where Q_i is defined in Eq. (73). **B)** Probability of detecting an FP as a function of x , calculated (yellow curve) and Monte Carlo simulated (blue dots).

along the x -axis. For each initial condition, we evaluate the FP detection probability as the ratio between the number of Monte Carlo runs where $X^{\text{diff}}(\Delta t) > X^{\text{diff}}(0) + \theta^*$ and the total number of Monte Carlo runs (Fig. 22B, blue dots), the result of which precisely fit with the calculated $\alpha(x)$. Finally, we obtain the overall FP detection probability, Γ_A , from Eq. (52), which, in this pure diffusion test case, differs from Γ_C by only 0.06%.

3.3.2 Validation of the iterative scheme

Before applying the iterative procedure (Fig. 20) to real data, we first validate it against time series generated by numerically integrating Eq. (47) (using the Euler–Maruyama scheme). We consider two validation cases: in Case 1 the amplitude of the jumps are comparable to the diffusive fluctuations and jumps are sparser in time (i.e., occur at a lower rate) than in Case 2, where jumps are much larger than the background noise and their rate is double that of Case 1 (Fig. 23A). The specific functions and parameters used to generate and analyze these validation data are shown in Table 3, which can be summarized as: low rate, low amplitude, high noise for Case 1, and high rate, high amplitude, low noise for Case 2. Preliminary tests with a linear drift function showed a successful fit between the effective SDE and the numerical data. We now opt for a more general and arbitrary shape where the drift function is nonlinear and non-monotonic.

The only restrictions are that it yields a single stable fixed point, and that the resulting stochastic process be stationary. Although the parameters and functions used for the simulations are known, they are not used in the inference procedure, only $\{X(t)\}$ is. To assess the performance of the proposed method, we compare the estimated \hat{D} , $\hat{\lambda}$, \hat{F} , and \hat{Q}_B with their true values.

Table 3: Parameters and functions for the validation cases

	Case 1	Case 2	Pure diffusion
λ	0.1	0.2	0
D	0.13	0.05	0.15
θ^*	0.125	0.07	0.1
μ	-1.2	1	N/A
σ	0.2	0.5	N/A
Q_B	Lognormal(μ, σ^2)		N/A
a	0.2		
F	$-(a(x - 0.5)^3 + 0.5a(x - 0.7)^2 + 0.1)$		
Δt (s)	10^{-4}		

We choose these two specific cases because they challenge both the sensitivity and the robustness of our method. In Case 1, the statistics of $\{X(t)\}$ are not too different from those of a pure diffusion process, which makes it easier to estimate F , but difficult to extract Q_B amidst the FPs. In contrast, the jumps in Case 2 are well separated from the diffusive fluctuations, which allows for a more direct estimation of Q_B . The presence of large jumps in this case, however, significantly alters $\{X(t)\}$ and its PDF, making it harder to estimate F . In both cases, however, we find that the original SDEs can be precisely recovered by our method. For instance, simulating Eq. (47) with \hat{D} , $\hat{\lambda}$, \hat{F} , and \hat{Q}_B of the last iteration not only produces time series that resemble the originals (Fig. 23A and 23B), but also yields an excellent fit between the reconstructed and original PDFs (Fig. 23C), with an $O(10^{-4})$ root-mean-square error (normalized by the range of $\{X(t)\}$), and ACFs (Fig. 23D).

Inspecting the results from the last iteration, we indeed see that the correct drift function is recovered (Fig. 24A). This is done through the use of the so-called differential Chapman-Kolmogorov equation. This is then used to calculate the next Q_A and Γ_A (Fig. 24B), which allows the correct Q_B to be demixed from the measured Q_C (Fig.

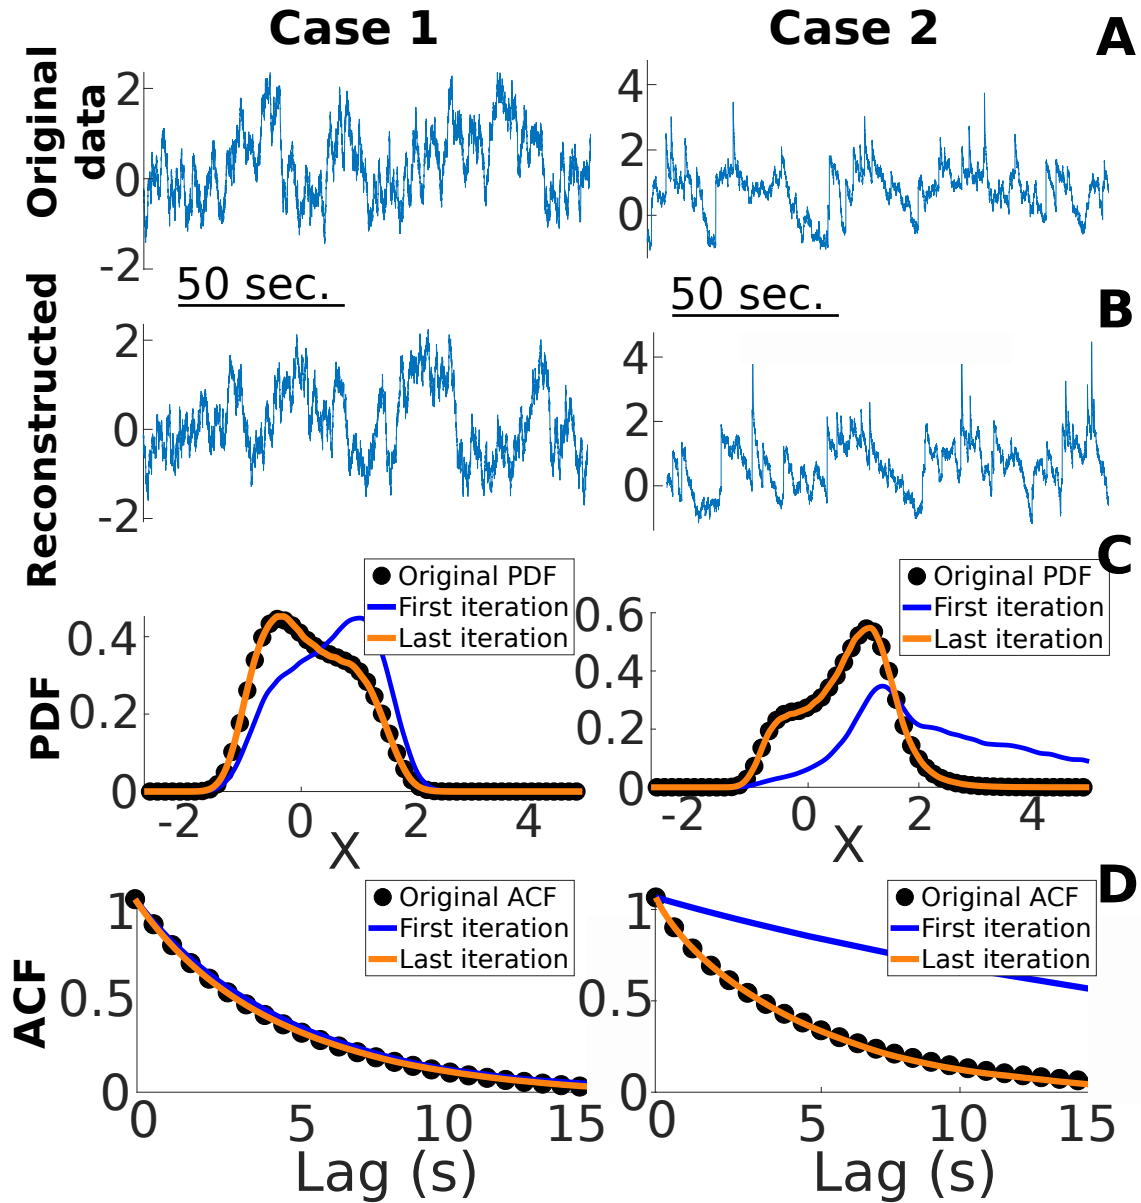


Figure 23: *Our inference method successfully yields an effective stochastic model of the original data. A)* Realizations of the jump-diffusion processes for the two different validation cases. **B)** Simulation results from the jump-diffusion SDE inferred by our method. **C)** Comparison between the PDF of the original validation data (black dots) and that of the last iteration of the effective SDE (yellow curve). The PDF associated with the first iteration is also shown (blue curve). **D)** Similar comparison between the original and estimated ACF.

24C). We also obtain low errors when comparing our estimates and the correct values of D and λ (Table 4). Note also that, although not shown here, we obtain the same fit quality between original and reconstructed when we apply our method to hybrid cases, where, for instance, the small jump amplitudes are paired with a higher rate instead of a lower one and vice-versa.

Table 4: Comparison between estimated and correct parameters.

Case 1				Case 2			
	Correct	Estimated	Error %		Correct	Estimated	Error %
D	0.13	0.13003	0.023	D	0.0500	0.05005	0.1
λ	0.1	0.0978	2.2	λ	0.2000	0.196	2.13

3.3.3 Application to experimental data

We now proceed with an application of our inference method to electrophysiological data published in Ref. [124]. These data consist of in vitro, intracellular recordings of membrane voltage fluctuations in pyramidal neurons of the weakly electric fish *Apteronotus leptorhynchus*. These fish are endowed with an active sensing mechanism whereby they generate a high frequency (~ 700 to 1000 Hz) oscillatory electric field around their body. This electric signal, along with any distortions caused by objects, preys, and conspecifics, is sensed by electroreceptors located on the fish’s body. This information is then sent to the hindbrain, where it reaches the first stage of electrosensory processing, called the electrosensory line lobe (ELL). The recordings in Ref. [124] were taken from neurons of the ELL, specifically in the centrolateral and centromedial segments (CLS and CMS, respectively). In order to isolate the impact of voltage-gated ion channels on membrane potential fluctuations, the ELL was treated with pharmacological agents (CNQX and APV) that block synaptic transmission onto the pyramidal cells. The resulting fluctuations are thus fully attributed to cell-intrinsic sources, which we refer to as membrane noise. The main source of this type of noise is often assumed to be the stochastic opening and closing of ion channels, i.e., channel noise [201]. In this case, we cannot rule-out other potential contributions, as non-trivial soma-dendrite interactions have been observed in these cells [146]. Note also that, by imposing different holding currents on the cells, Ref. [124] recorded ongoing membrane noise at various levels of hyperpolarization relative to

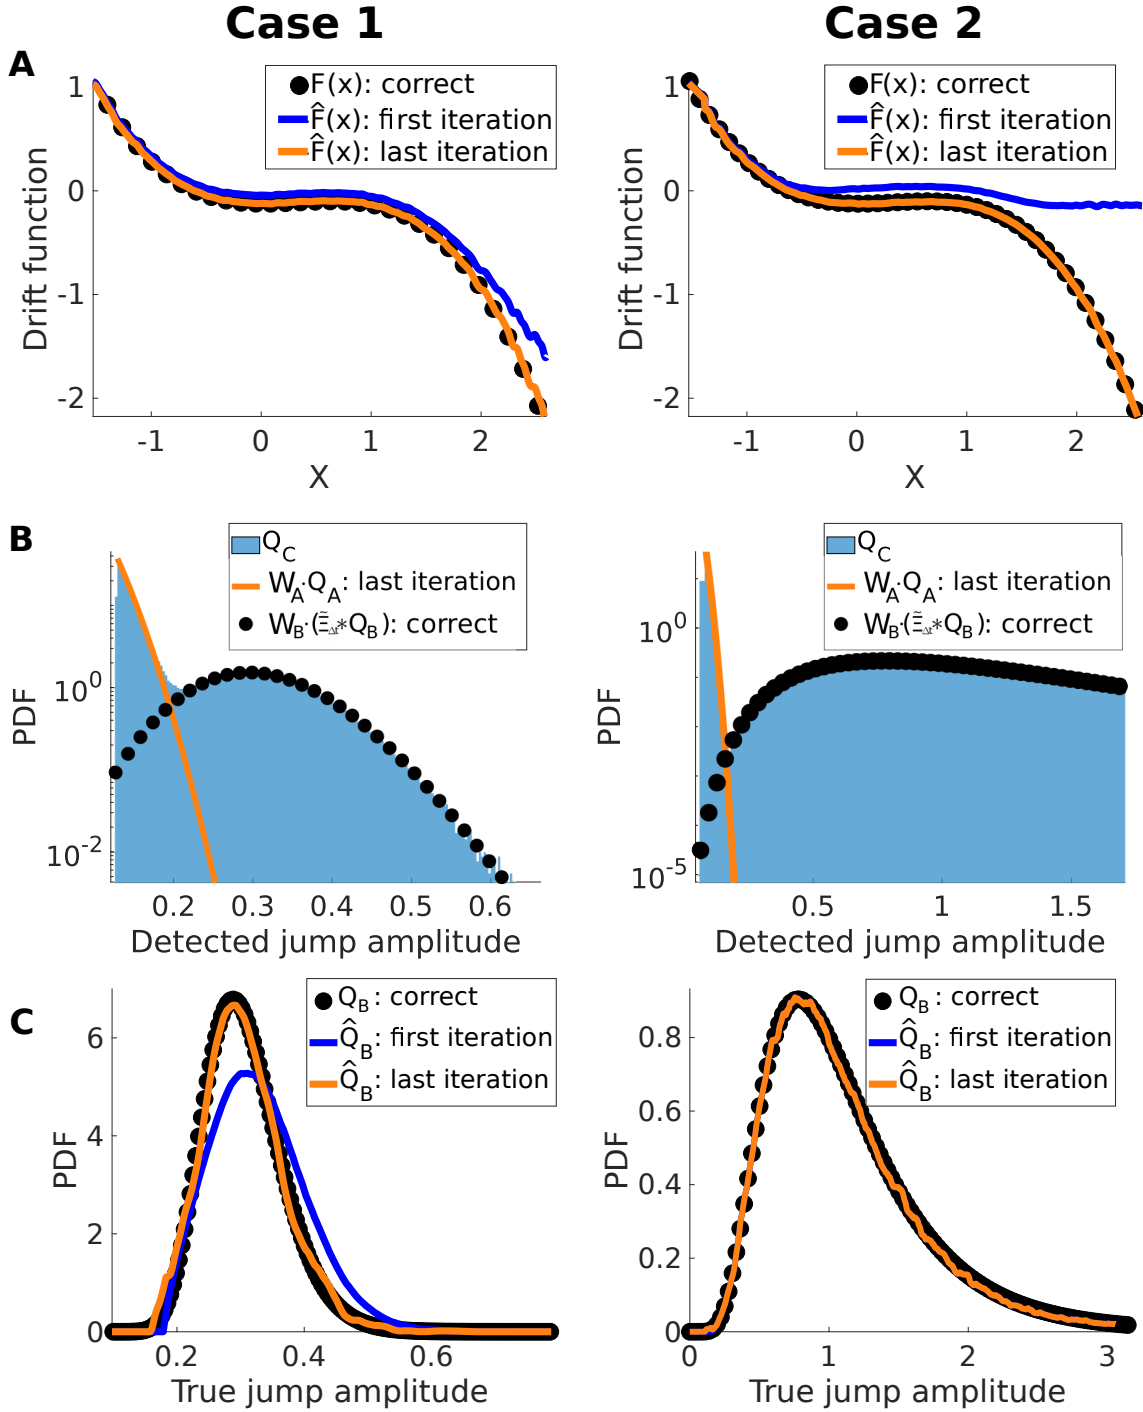


Figure 24: Our inference procedure recovers the correct parameters and functions from Eq. (47). **A)** The correct drift function (black dots) is recovered in the last iteration (yellow curve). First iteration results are shown for comparison (blue curves). **B)** We see that the measured amplitude distribution, Q_C , is in fact a mixture of Q_A (yellow curve), and $\tilde{\Xi}_{\Delta t} * Q_B$ (black dots, see Eq. (66)). **C)** After deconvolving the latter, we do recover the correct Q_B (yellow curve). Note that in Case 2, the first and last iteration are confounded, as the true jump amplitudes are almost directly separable from those of FPs.

spike threshold.

Of interest here is the presence of large, jump-like events, called blips, that abruptly depolarize the cells (Fig. 25A, asterisks). Although Ref. [124] puts forth a hypothesis as to the functional role for these blips, the mechanism underlying their occurrence is unknown. This, along with the limited amount of data, hinders the development of any meaningful mechanistic model of this phenomenon. The effective jump-diffusion approach developed here, however, is particularly well-suited to circumvent this knowledge gap. Indeed, the resulting phenomenological models provide a useful tool for dynamically interpreting the available data without relying on poorly constrained biophysical mechanisms. For instance, we can address questions such as: does certain parameters or functions of the model change as a function of the mean membrane potential?

Here we analyze recordings from two CMS and two CLS cells, each with 5 or 6 levels of hyperpolarization: from -25 to 0 mV below threshold, with 5 mV steps between levels. Using these relative levels with respect to spike threshold is required to compare cells that might have different thresholds (e.g., -67 to -63 mV for CMS cells [124]). After removing experimental artifacts (see Sect. 3.3.4), we obtain a total of 23 traces, each lasting approximately 10 s. Applying our inference method to these traces yields a good fit between the resulting simulations and the original data (Fig. 25A): the PDFs differ only by $O(10^{-2})$ normalized root-mean-square errors, and the power spectra fall within 95% confidence intervals of each other (Fig. 25B and C).

Further insight can be gained by comparing the estimated SDE parameters and functions, \hat{D} , $\hat{\lambda}$, \hat{Q}_B and \hat{F} , across all traces. We thus see, for instance, that CLS cells increase their jump rate (Fig. 26A), but not jump amplitudes (Fig. 26B), as they approach threshold. Note also that we could not measure any significant jump component for CMS cells. Instead, fluctuations in these cells are well-described by pure diffusion. Furthermore, all cells show an increase in their diffusive noise intensity with depolarization, and this is more prominent in CMS cells (Fig. 26C). Lastly, to compare the different drift functions with a scalar measure, we apply a linear fit to $F(x)$ (estimated as in the previous section) in the vicinity (± 0.2 mV) of the stable fixed point. The slope parameter resulting from this fit can be interpreted as a measure of how wide or narrow the potential function is around the resting membrane voltage. Using this measure, we

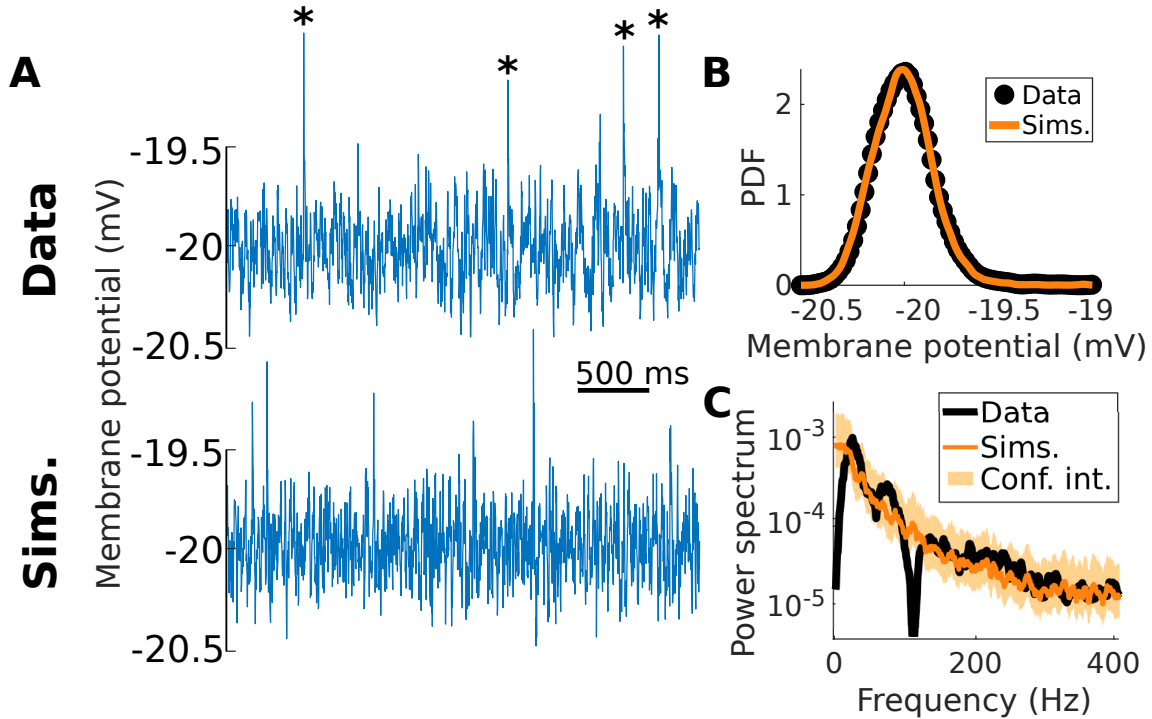


Figure 25: *Membrane noise in CLS cells can be modelled with an effective jump-diffusion SDE. A)* An exemplar recording of membrane voltage fluctuations in a CLS pyramidal neuron (top). A holding current is used to maintain the cell at 20 mV below its spike threshold. Asterisks show the four largest blips found in this trace. Under the assumptions listed in Sect. 3.2.1, simulations of the effective SDE (bottom) are qualitatively similar to the original data (here $\Delta t = 1.2$ ms). **B)** This similarity is confirmed by the close match between the data and simulation PDFs (yellow curve and black dots, respectively). **C)** Despite having no role to play in the inference procedure, the power spectrum of the data (black) also fits with that of the simulations (yellow). The notches in the power spectrum of the data result from the removal of experimental artifacts.

find no systematic intra-cell trend, but we do observe large differences between cell types: CLS cells have a wider potential function than CMS ones (Fig. 26D).

3.3.4 Data processing

For each cell, the raw data consist of a continuous, 60 to 70 second staircase-like trace, sampled at 20 kHz. Each step lasts ~ 10 s and corresponds to a different holding current, which was applied such as to create 5 mV hyperpolarization from the previous level. In order to segment the recordings into different traces for each level, we first identify the transition times between different holding currents. This is done visually, as the transitions are unambiguous, and we omit ± 0.5 seconds around those times. At the -5 and 0

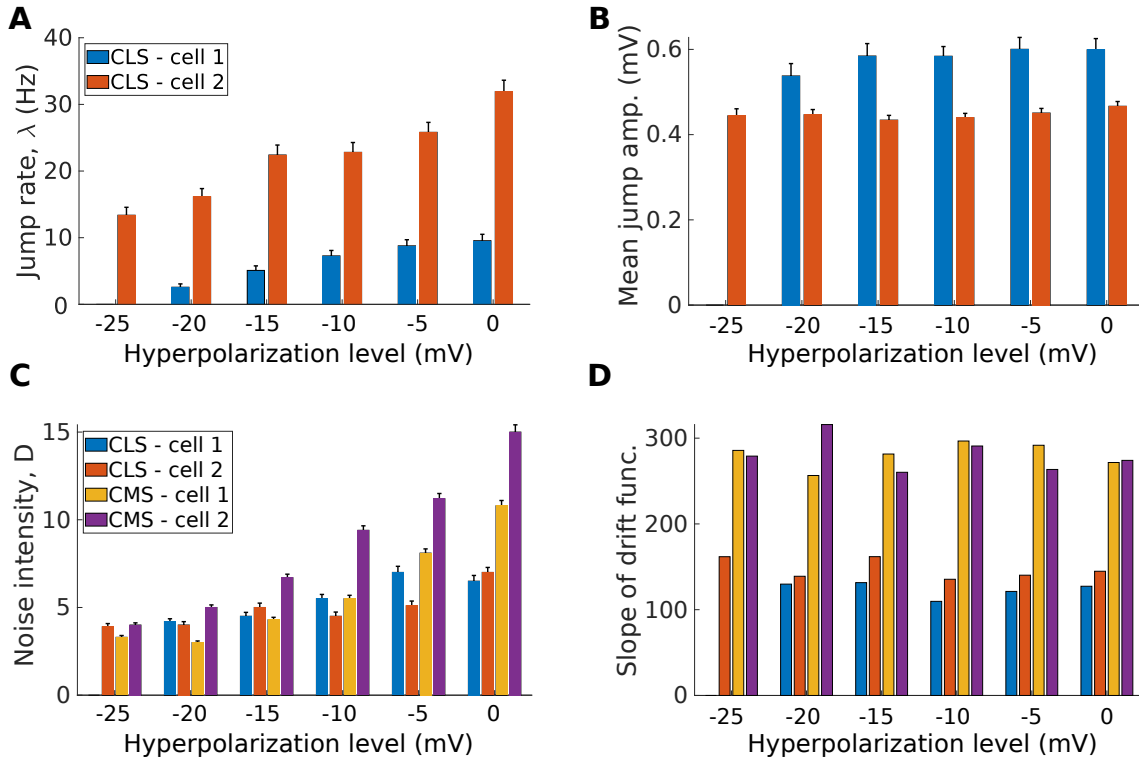


Figure 26: *CLS cells increase their rate and noise intensity, but not jump amplitudes, when they approach threshold, and all cells maintain a steady drift function across levels.* **A)** Jump rate of CLS cells. Error bars show one standard deviation, assuming that the number of detected blips is Poisson. **B)** Mean amplitudes of the blips. Error bars show one standard deviation, calculated from 1000 bootstrap samples of the original amplitude values. We observe a similar lack of systematic trend in the variance of the amplitudes (not shown). **C)** Noise intensity for all cells. Error bars show one standard deviation, calculated from 1000 bootstrap samples of the original data increments. **D)** Local slope of the drift function, as determined from a linear fit over a ± 2 mV range around the stable fixed point. Error bars are too small to see on this scale, but are calculated as 95% confidence interval of the slope parameter of the linear least square fit.

mV level, a few spikes (1 to 4) occur in the recordings. They are manually removed from the traces along with the ensuing refractory period.

A conspicuous aspect of the resulting traces is the presence of slow, large amplitude (1 Hz, ~ 1 mV) quasi-oscillations overlaid with the faster, stochastic fluctuations. The exact source of this slow component is unknown, but is possibly related to persistent sodium channels, which have been shown to populate the soma and proximal apical dendrites of ELL cells [124], and to produce slow perithreshold oscillations in entorhinal stellate neurons [40]. In any case, these slow oscillations are outside the scope of the method

presented here. A moving average filter (0.05 s window size) is thus applied to remove this low frequency content from the signal.

Line noise is removed at all 60 Hz multiples with a notch filter, but the data are also contaminated with artifacts in other frequency bands, potentially from interference with other sources. This is most prominent in the 900 - 3000 Hz, but also carries around lower frequencies, e.g., 100, 270, and 550 Hz. To account for this artifact, we opt for the combination of a low-pass filter with a 900 Hz cutoff, and 20 Hz wide band-stop filters centred on the other problematic frequencies. Electrophysiological recordings can be dominated by measurement noise at high frequencies [39]. In our case this is seen as a flattening out of the PSD above 1000 Hz, so the 900 Hz cutoff used here does not lead to the loss of important biological signals.

The end result of this processing chain are time series that exhibit fluctuations typical of jump-diffusion processes. We do observe, however, significant higher-order correlations on the smallest timescales ($O(\Delta t)$, $\Delta t = 50 \mu\text{s}$). To quantify these correlations, we use the notion of the Einstein-Markov timescale [54]. This is a measure of the timescale below which the Markov property no longer holds. Stochastic time series often show a departure from the Markov property on small timescales, possibly due to noise source correlation, the presence of an inertial component in the dynamics, or measurement noise [54]. Following [54] and [6], we estimate this Einstein-Markov timescale by finding the value of τ that minimizes

$$\chi^2 = \int \int \int \frac{[P(x_1, x_2, x_3) - P(x_3|x_2)P(x_2, x_1)]^2}{\sigma^2} dx_1 dx_2 dx_3, \quad (75)$$

where $x_1 = x(t)$, $x_2 = x(t + \tau)$, $x_3 = x(t + 2\tau)$, and σ^2 is sum of the traces of the covariance matrices associated with the distributions in the numerator. For a proper Markov process, $\chi^2 = 0$, $\forall \tau$. In this case, we find the minimum of χ^2 at 1.2 ms, indicating that the Einstein-Markov timescale of the data is over one order of magnitude larger than the sampling interval, Δt . This means that, on the time scale of individual observations, the data evolve with a history dependence that is incompatible with a Markovian description. If, however, we look at the data on a coarser time scale, e.g., the Markov-Einstein time scale of 1.2 ms, then the Markov property is approximately satisfied. In that case only can we hope to use Eq. (47) as a valid model for these data. To account for this problem, we resample the data at a 1.2 ms interval (~ 830 Hz sampling rate), and obtain

the final time series on which to apply our method (Fig. 25A,top), with the time step equal to the Markov-Einstein time scale. Note that this situation is conceptually similar to how the Langevin model of diffusion (where the position of a particle is, by itself, not Markovian) reduces to the Einstein model (where the position is Markovian) only above a certain time scale [62].

3.4 Discussion

In this study, we develop an iterative procedure that recovers the parameters and functions of a jump-diffusion SDE, based solely on a realization of the associated stochastic process. This approach is validated when the jumps are comparable in size to the diffusive fluctuations (Case 1), as well as when they are much larger than diffusive fluctuations (Case 2). We apply this method to membrane voltage fluctuations recorded in pyramidal neurons of electric fish. Our analysis reveals that these data can indeed be represented as jump-diffusion processes. We find that pyramidal neurons increase their jump rate and noise intensity as they approach spike threshold, while their jump amplitudes and drift function remain unchanged.

Our approach relies on five main components: the use of the differential Chapman-Kolmogorov equation to estimate the drift function, the use of quadratic variation on jump-free segments to estimate the diffusive noise intensity, the detection of jumps via threshold-crossing of the increments, the modelling of detected jumps as a mixture of true jumps and FPs, and the calculation of FP statistics used to extract true jump statistics from the detected jump pool.

Although we estimate the drift function and the true jump amplitude distribution non-parametrically, we do limit our study to the case of additive diffusive noise, of constant jump rate, and of Poisson jumps. Relaxing the additive noise assumption would require an estimation scheme for the diffusion function, $D(x)$. For purely diffusive processes, this function can be obtained directly through the estimation of the second Kramers-Moyal coefficient, which is defined in terms of the second conditional moment of the increments. Evaluating this moment simply requires the knowledge of the conditional PDF across time steps. For a Poisson jump-diffusion process, however, Ref. [3] has shown that the diffusion function can in fact be expressed in terms of the second

conditional moment of the increments, the jump rate, and the second moment of the jump amplitudes. It should thus be possible to include the estimation of $D(x)$ into the iterative portion of our method (Fig. 20). Indeed, estimates of the jump rate and of the amplitude distributions could be used at each pass to estimate the diffusion function.

As for the assumption of constant jump rate, it should be possible to extract a rate function, $\lambda(x)$, as long as an x -dependent version of Eq. (64) can be written. This would require a long enough data time series such as to produce an estimate of $\Gamma_C(x)$. Relaxing the assumption of Poisson jumps, however, would be more difficult to do. The detection probability of true jumps, $\Gamma_B = \lambda\Delta t$, would obviously need to be modified with the appropriate expression. Moreover, the specific form of the differential Chapman-Kolmogorov used here, Eq. (71), relies on the assumption of Poisson jumps (see Appendix A) and would thus need to be extended in a manner that depends on the precise non-Poissonian nature of the jump process. More specifically, the last two terms in Eq. (71), which are originally defined based on the transition rates of the Poisson jump process, would now be derived from the modified Γ_B . Note that, for the special case of true jumps with zero-mean amplitudes, the drift function can be estimated directly from the first conditional moment of the increments, without relying on Eq. (71) [3].

3.4.1 Membrane noise

The unusual characteristics of membrane noise observed in CLS neuron, initially reported in Ref. [124], and represented here as jump-diffusion SDEs, might be implicated in novel ways in electrosensory processing. The analysis we perform here is a first step toward investigating this possibility computationally.

The positive impact of noise on information processing in neural systems is widely recognized [129]. Although channel noise was initially thought to be too weak compared to synaptic noise to influence a neuron’s output statistics [169, 21], it has since been shown to significantly impact neuronal reliability and action potential timing [46, 170, 47, 28, 40, 65]. Because it arises from the stochastic opening of voltage-gated ion channels, channel noise has been successfully modelled by populations of Markov chains with voltage-dependent transition rates [201]. In the quest for more computationally efficient models, however, various approximations have been used to model the

collective behaviour of these Markov chains as simple diffusion SDEs [52, 53, 66, 65] (sometimes called the diffusion approximation for channel noise [118, 147], in reference to the diffusion approximation⁸ for synaptic bombardment). One of these approaches, for instance, introduces a current noise term directly in the membrane equation [166, 65], effectively modelling the subthreshold voltage locally as an Ornstein–Uhlenbeck process. It is perhaps not surprising then that we obtain a good match between the observed CMS membrane noise (at various holding potential) and a pure diffusion SDE. In cases where membrane noise is more accurately described by multiplicative conductance noise [65], however, we might expect deviation from the simple SDE used here, similar as to how the diffusion approximation can misrepresent the subthreshold voltage distribution for certain types of synaptic drive [163].

We cannot completely exclude the possibility that small, hard to detect blips occur in CMS cells. Although our method proves capable handling this type of situation (Sect. 3.3.2), the limited amount of available experimental data in this case precludes us from conclusively ruling out the existence of a jump component in models of CMS membrane noise. In addition, in both type of cells we find a positive correlation between the noise intensity and the holding potential (Fig. 26). This is consistent with how membrane potential variance has been observed to increase with depolarization in these same cells [124], as well as in rat neocortical pyramidal neurons [89]. Simple Markov models involving only Na^+ and K^+ channels are able to reproduce this correlation [184].

The presence of blips in CLS neurons suggest that ion channels co-activate to produce abrupt depolarizing currents. It was shown in [194] that sodium channels appear in clusters, or hot spots, in ELL pyramidal cells. This might allow local depolarization of the membrane to sufficiently couple channels within a cluster. Alternatively, perhaps channels are physically and functionally coupled through scaffolding protein complexes, as observed in [209]. Regardless of the exact coupling mechanism, the commonly used assumption of independence between channels [201, 65] is likely violated by these blips. The fact that we have successfully constructed an effective jump-diffusion SDE for CLS membrane noise suggest that a diffusion-like approximation could be applied in this case as well. Such a deductive approach would, however, require tentative descriptions of local channel coupling to be included in the kinetic schemes.

⁸In this case we refer to how stochastic synaptic currents, modelled as shot noise and assumed to arrive at a high rate, can be approximated as a diffusion process.

Given the unknown biophysical mechanism underlying the blips, our effective jump-diffusion model is uniquely positioned to address questions related to their functional role. Future work will thus aim to incorporate our effective jump-diffusion model as a membrane noise term in a more complete model of CLS cells [18]. Recall, however, that we have obtained in Sect. 3.3.3 several inferred jump-diffusion SDEs, each corresponding to a specific level of depolarization. In order to produce a model of the full dynamics of CLS cells (over a wide range of depolarization levels), it will thus be necessary to combine these SDEs in a way that captures how the noise intensity, blip rate, and blip amplitudes depend on the membrane potential (Fig. 26). This could be done phenomenologically by considering a single SDE with a voltage-dependent noise intensity and blip rate. Furthermore, by accounting for the synaptic input associated with electrosensory input, the resulting model could investigate the possibility that blips assist or influence spiking, perhaps through a stochastic resonance-like phenomenon. Stochastic resonance, and more generally, stochastic facilitation, has been shown to be mediated by channel noise in models of auditory brain stem neurons [169] and in modelled neuronal arrays [4], as well as to be mediated by synaptic noise in models of neocortical pyramidal neurons [168]. Since blips share a similar shape and amplitude as AMPA-driven excitatory post-synaptic potential, this begs the question of whether stochastic resonance is at play in the detection of weak electrosensory signals, such as small prey. This was indeed hypothesized, but not explicitly shown, in Ref. [124]: the voltage-dependence of membrane noise makes it impossible to vary the noise level, the rate of occurrence, or the amplitude of blips, independently of the membrane potential. The value of this hypothesis, however, could be assessed by performing a numerical experiment with our effective SDE representation of blip-laden membrane noise.

Lastly, although we apply our method here to recordings where synaptic input is completely blocked, it might be applicable to certain type of synaptic input patterns, such as correlated bombardment. We also hope to apply this method to fluctuations of the active-sensing rate of pulse-type electric fish, where jump-like events also occur [97].

3.5 Appendix A

The particular form of the differential Chapman-Kolmogorov equation that we use here, Eq. (71), relies on the assumption of Poisson jumps. For completeness, and to make explicit the dependence of this equation on Poisson statistics, we provide the following

short derivation:

Let $X(t)$ be a continuous-time Markov process, possibly with discontinuous sample paths, in one dimension and with transition probability $p(x, t|x_0, t_0)$. The Chapman-Kolmogorov equation then has the following differential form [57]:

$$\begin{aligned} \frac{\partial p(x, t|x_0, t_0)}{\partial t} = & -\frac{\partial}{\partial x} [F(x, t)p(x, t|x_0, t_0)] + \frac{\partial^2}{\partial x^2} [D(x, t)p(x, t|x_0, t_0)] \\ & + \int [W(x|\tilde{x}, t)p(\tilde{x}, t|x_0, t_0) - W(\tilde{x}|x, t)p(x, t|x_0, t_0)] d\tilde{x}, \end{aligned} \quad (76)$$

where

$$W(x|\tilde{x}, t) \equiv \lim_{\Delta t \rightarrow 0} \frac{p(x, t + \Delta t|\tilde{x}, t)}{\Delta t}. \quad (77)$$

The first and second term of Eq. (76) are associated with the continuous part of $X(t)$ and are often referred to as the advective and diffusive terms in the parlance of partial differential equations [150]. The last term in Eq. (76), on the other hand, represents how probability mass is transferred through discontinuities in $X(t)$, as the function W is the instantaneous transition rate between different values of x . Indeed, sample path continuity is expressed as $W(x|\tilde{x}, t) = 0$ [57], and in that case, Eq. (76) reduces to the Fokker-Planck equation. Note that, in the text, we apply the commonly used practice of abbreviating the transition probability as simply “the PDF of $X(t)$ ” [62]. Also note that, in the case where both the drift and diffusion functions vanish, Eq. (76) reduces to the so-called Master equation in statistical mechanics [61, 151].

In order to evaluate Eq. (77), we first introduce three mutually exclusive and collectively exhaustive events: $E_0 = \{\text{no jumps in } [t, t + \Delta t]\}$, $E_1 = \{\text{exactly one jump in } [t, t + \Delta t]\}$, and $E_2 = \{\text{two or more jumps in } [t, t + \Delta t]\}$, such that $E_i \cap E_j = \emptyset$, for $i \neq j$, and $\text{Prob}(\cup_i E_i) = 1$. From the law of total probability, we get:

$$p(x, t + \Delta t|\tilde{x}, t) = \sum_{i=1}^3 p(x, t + \Delta t|\tilde{x}, t; E_i) \cdot \text{Prob}(E_i). \quad (78)$$

Assuming that discontinuities occur through a Poisson process with parameter λ , we can

use the well known results [149, 85]:

$$\begin{aligned}
\text{Prob}(E_0) &= 1 - \lambda\Delta t + o(\Delta t) \\
\text{Prob}(E_1) &= \lambda\Delta t + o(\Delta t) \\
\text{Prob}(E_2) &= o(\Delta t),
\end{aligned} \tag{79}$$

where $o(\Delta t)$ goes to zero faster than Δt :

$$\lim_{\Delta t \rightarrow 0} \frac{o(\Delta t)}{\Delta t} = 0. \tag{80}$$

Inserting the sum in Eq. (78) into the limit of Eq. (77) results in the sum of three limits, associated with E_0 , E_1 , and E_2 :

$$W(x|\tilde{x}, t) \equiv \lim_{\Delta t \rightarrow 0} \sum_{i=0}^3 \frac{p(x, t + \Delta t|\tilde{x}, t; E_i) \cdot \text{Prob}(E_i)}{\Delta t}, \tag{81}$$

where the first term vanishes because of sample path continuity between jump events: $\lim_{\Delta t \rightarrow 0} p(x, t + \Delta t|\tilde{x}, t; E_0)/\Delta t = 0$, while the third one vanishes as per Eq. (80): $\lim_{\Delta t \rightarrow 0} \text{Prob}(E_2)/\Delta t = 0$. If jump amplitudes are distributed according to Q_B , then this distribution can be defined as $Q_B \equiv \lim_{\Delta t \rightarrow 0} p(x, t + \Delta t|\tilde{x}, t; E_1)$, which results in:

$$\begin{aligned}
W(x|\tilde{x}, t) &= \lim_{\Delta t \rightarrow 0} \frac{p(x, t + \Delta t|\tilde{x}, t; E_1) \cdot \text{Prob}(E_1)}{\Delta t} \\
&= \lim_{\Delta t \rightarrow 0} \frac{p(x, t + \Delta t|\tilde{x}, t; E_1) \cdot (\lambda\Delta t + o(\Delta t))}{\Delta t} \\
&= \lambda Q_B(x - \tilde{x}).
\end{aligned} \tag{82}$$

When inserted into Eq. (76), this yields

$$\begin{aligned}
\frac{\partial p(x, t|x_0, t_0)}{\partial t} &= -\frac{\partial}{\partial x} [F(x, t)p(x, t|x_0, t_0)] + \frac{\partial^2}{\partial x^2} [D(x, t)p(x, t|x_0, t_0)] \\
&\quad - \lambda p(x, t|x_0, t_0) + \lambda \int Q_B(x - \tilde{x})p(\tilde{x}, t|x_0, t_0)d\tilde{x}.
\end{aligned} \tag{83}$$

If we assume a time-homogeneous drift function, additive noise, and a positive support for Q_B , we recover Eq. (71).

Transitions between stochastic and oscillatory active sensing in pulse-type electric fish

4.1 Introduction

Most senses are active [172, 197, 156]. Rather than wait passively for signals to be detected by their sensors, animals actively perform motor routines (e.g., whisking, sniffing, touching, eye movements) or emit energy in the medium (e.g., echolocation, electrolocation) in order to purposefully gather information from their environment. This information is directly linked to an animal's movements, which are instrumental in shaping the spatiotemporal structure of the sensory flow [80, 82]. Furthermore, when these movements are performed rhythmically, reafferent⁹ sensory streams are also rhythmic, the temporal regularity of which is advantageous for sensory processing [172]. A large body of evidence indeed shows that ongoing, low-frequency neuronal oscillation in sensory processing areas are entrained to rhythmic stimuli [68]. This allows the high-excitability

⁹Sensory input resulting from an animal's own movement.

phase of neuronal oscillations to coincide with the arrival of sensory input, which provides a beneficial modulation of sensory gain [111]. This bottom-up entrainment can also be modulated by top-down¹⁰, task-specific attentional signals, leading to the hypothesis that entrainment of low-frequency neuronal oscillations is a core mechanism of selective attention [109, 171, 172, 22]. The link between active sensing, motor rhythms, and attention is, however, only beginning to be understood. There is thus a need for robust, observable, and controllable manifestations of this link across different active sensory modalities.

Here, we report on and characterize an hitherto unknown behavioural state of pulse-type weakly electric fish during which the rate of electrosensory acquisition becomes rhythmic and coupled to low-frequency movement around a specific stimuli. The oscillatory dynamics of this sampling strategy is in stark contrast to that exhibited during bouts of exploratory movement or during rest, both of which display instead stochastic fluctuations of sensory sampling. We further show that these fluctuations are well-modelled by jump-diffusion stochastic processes, and that the qualitative differences between rest and exploratory electrosensory sampling can be explained in dynamical terms.

Similarly to echolocating animals, pulse-type electric fish sense their environment through a set discrete sampling events known as electric organ discharges (EOD). Each EOD consists of a brief, bi-phasic, and dipolar electric field that permeates the surroundings. Perturbations to this field, caused by objects such as preys, are detected by an array of cutaneous electroreceptors distributed over the body of the fish (with a higher concentration around the head). By performing certain stereotyped movement patterns [191, 44, 81], and by voluntarily modulating the inter-EOD intervals [97], these fish exhibit a rich set of sensorimotor strategies [80].

Electrosensory acquisition in pulse-type fish presents several advantages as a target for active sensing research: i) individual EODs can be non-invasively detected and provide a precise time-stamp, across all receptors on the body, for the arrival of sensory input (there are negligible propagation delays in contrast to auditory signals), ii) electroreceptors are fixed on the body and, therefore, move through space only from body

¹⁰Emanating from a brain area, or a network “higher up” in the information processing hierarchy, where the lowest level corresponds to most direct and detailed encoding of sensory stimuli, while higher levels correspond to more integrated information that can involve cognitive processes.

movements (as opposed to, e.g., whisking and eye movements, where the sensors themselves move in superposition to head and body movements), iii) EODs are stereotyped (as opposed to bats, which can modulate the duration of their calls [93] and change the width of the sound beam [107]), and iv) fish cannot dynamically alter the properties of their electroreceptors (as opposed to how bats can alter the shape of their ears to modify incoming signals [56]). Sensorimotor integration in pulse-type fish thus relies entirely on a combination of body movement and temporal modulation of the inter-EOD interval, henceforth called inter-pulse interval, or IPI, both of which are readily measurable in a laboratory setting.

In Sec. 4.2, we first present our effective stochastic model for the IPI fluctuations during rest and exploratory movements, then we describe the phenomenology of oscillatory sensory sampling epochs. Next, we discuss in Sec. 4.3 the possible implications of this phenomena in the context of attention research. Details pertaining to data analysis and modelling are presented in Sec. 4.4.

4.2 Results

We present here our analysis of a previously published dataset [95], in which pulse-type electric fish *Gymnotus sp.* were monitored during long periods of time (~10 hours) while behaving freely in a circular tank. The thermoregulated tank was surrounded by a sensory isolation chamber that blocked all external lights, sounds, and vibrations. As fish were unperturbed for the duration of the recordings, the observed behavioural states, and any transition between them, were thus internally and spontaneously generated. EODs were recorded from a set of electrodes placed around the periphery of the tank.

The resulting data contains information on the amplitude of the EODs and on when they occur. On one hand, information on fish movement is extracted from the EOD amplitude measurements. Indeed, fish movement causes temporal variation in the measured EOD amplitudes due to change in the relative position and orientation between the fish and the electrodes [96]. These variations are then quantified into a movement variable, which can be unambiguously thresholded to determine whether or not the fish is moving [94]. On the other hand, the EOD time-stamps are converted into a time series for the IPI.

Significant variability in the IPI time series is coupled to movement (Fig. 27A). For

instance, the IPIs are consistently shorter during movement than during rest, indicating that sensory samples are acquired at a greater rate during locomotion. The IPI distribution is in fact bimodal, the modes of which correspond to a moving state (shorter IPIs) and a resting state (longer IPIs). It has been shown elsewhere that the switching statistics between these two preferred values can be modelled by the stochastic dynamics of a double well system, where the relative depth of the wells is slowly modulated by an external forcing [132].

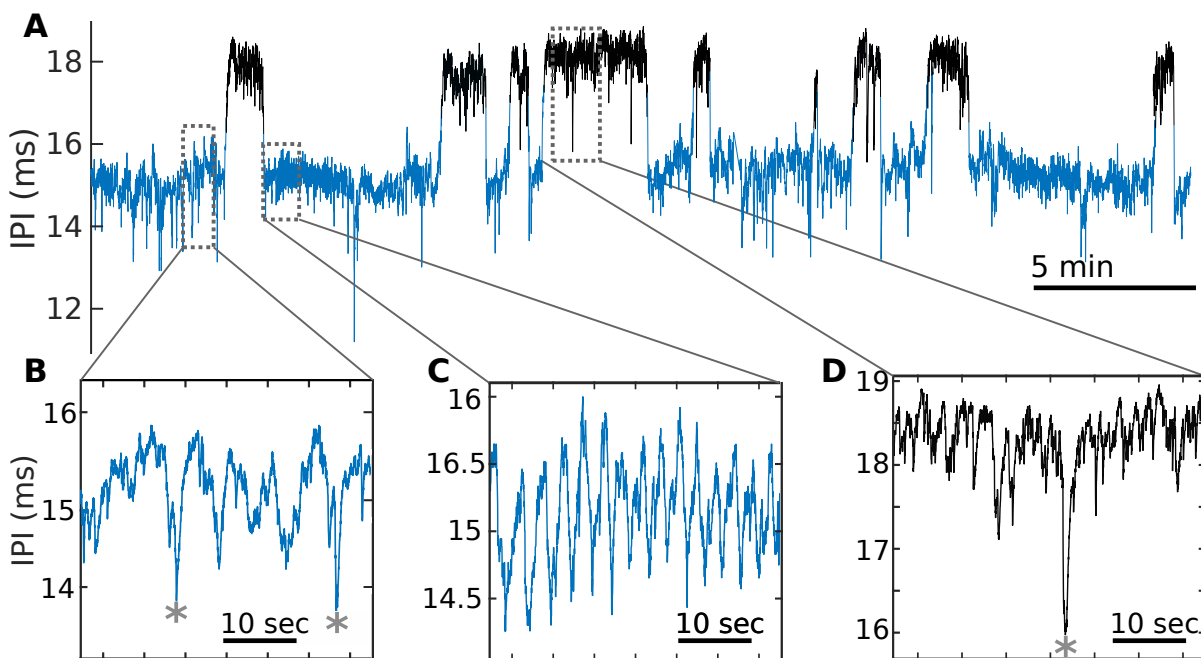


Figure 27: *Electrosensory sampling during movement alternates between distinct dynamical regimes.* **A)** IPI recording of a freely behaving electric fish. Fish sample their environment at shorter intervals during movement (blue traces) than during rest (black traces). **B) and C)** During movement, fish exhibit either stochastic fluctuations of the IPIs (**B**), or bouts of oscillatory sensory sampling (O-states, **C**). The former are populated with abrupt transients of rapid sensory sampling (E-scans, asterisks). **D)** During rest, IPIs evolve stochastically and are also populated with E-scans.

The dynamics of the IPI on shorter timescales, however, have not yet been examined and are far from trivial. For instance, during movement, fish alternate between two modes of electrosensory sampling: stochastic (Fig. 27B) and oscillatory (henceforth called O-states, Fig. 27C). Respectively, these are associated with exploratory locomotion (mostly along the walls of the tank), and with repeated side-to-side movements in front of the nylon locking nuts used to seal the gates between the tank and the housing

compartments. These O-states can last for up to four minutes. To our knowledge, such behaviour has not yet been documented and is thus a novel addition to the behavioural repertoire of *Gymnotus sp.*. During rest, the IPI also evolve stochastically, albeit with more stable fluctuations than during movement (Fig. 27D).

Furthermore, the IPI time series is conspicuously populated by short and abrupt transients where EODs occur in rapid succession, i.e. in a brief shortening of the IPIs (henceforth called E-scans [97], Fig. 27A and 27D, asterisk). E-scans occur at seemingly random times, and display a range of amplitudes. The E-scan rate has been shown to be modulated by reward expectation during food-searching tasks and is therefore believed to be involved in spatial attention [97].

In the following subsections, we perform data analysis to characterize the different modes of electrosensory acquisition, and we present low-dimensional stochastic models as a means to understand the core dynamics underlying these active sensing strategies.

4.2.1 Stochastic active sensing

Despite the lack of data pertaining to the neuronal origin of E-scans, we can still interpret the IPI fluctuations from a dynamical perspective. Indeed, we show here that the stochastic evolution of the IPIs, during both rest and exploratory movement, are well-modelled by the combination of diffusive fluctuations along with abrupt jumps of random amplitude. This is formalized by a jump-diffusion SDE of the form:

$$dX_t = F(X_t)dt + \sqrt{2D}dW_t + dJ_t, \quad (84)$$

where F is called the drift function, D the noise intensity, and W_t is a Wiener process (i.e., Brownian motion). The E-scans are represented by the events (referred to as jumps) of the compound Poisson process, J_t :

$$J_t = \sum_{i=0}^{N_t} A_i, \quad (85)$$

where N_t is a Poisson point process with rate λ , and the A_i 's are the independent and identically distributed jump amplitudes with distribution Q_A .

By applying the data-driven method developed Chapter 3, on stationary epochs of

the IPI time series, we extract semi-parametric estimates of D , λ , Q_A , and F . This results in an effective stochastic model of the original data. We obtain realizations of this model by integrating Eq. (84) numerically (Fig. 28A). Comparing the statistics of the numerical solutions to those of the original time series allows us to evaluate whether the effective model is indeed an adequate representation of the data.

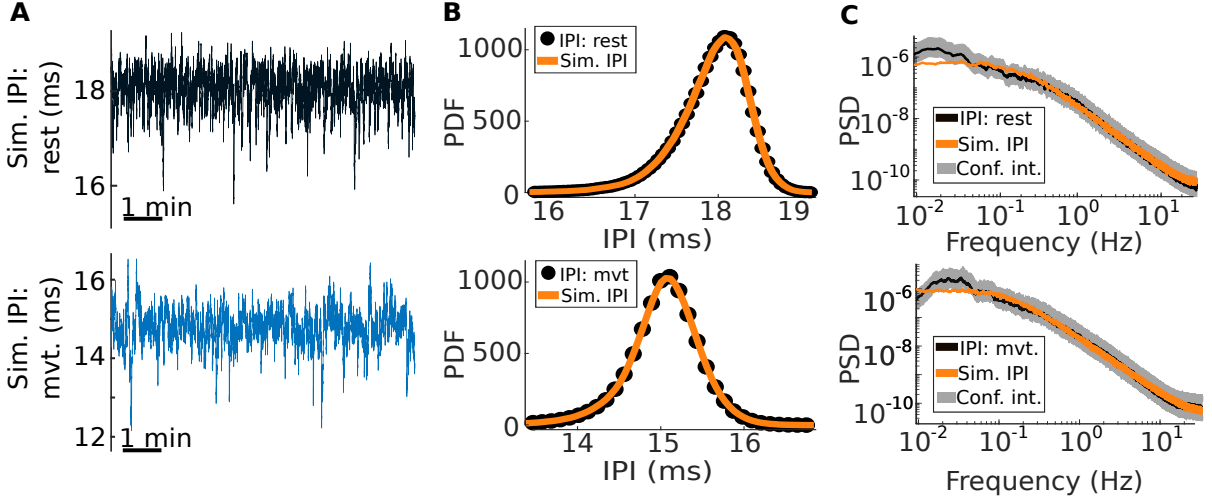


Figure 28: *Stochastic fluctuations of the IPI during rest and movement can be modelled by jump-diffusion SDEs. A)* Realizations of the effective stochastic models of IPI fluctuations during rest (top) and movement (bottom). **B)** Comparison between the PDF of the IPI time series (black dots) and that of the simulated models (yellow curves) during rest (top) and movement (bottom). **C)** Comparison between the PSD of the IPI times series (black curves) and that of the simulated models (yellow curves) during rest (top) and movement (bottom). The grey shaded area represents the 95% confidence interval.

We find that both the probability density function (PDF) and the power spectral density (PSD) of the IPIs closely match those of the simulations, during both rest and movement (Fig. 28B and C). Also note that, although PDFs of the original data are used during the estimation of the drift function, the PSDs are not in any way involved in the inference procedure.

Further insight can be gained by examining the differences between the effective models associated with rest and with exploratory movement. Notably, we find that the jump rate, λ , and diffusive noise intensity, D , are both higher during rest than during movement (Table 5). Furthermore, we find that the drift functions, F , are both non-linear, and that the one associated with rest is steeper (Fig. 29A), indicating a narrower potential well. Lastly, we find a slight but significant difference between the jump amplitude

distribution, Q_A , for rest vs. movement ($P < 10^{-3}$, two-sample Kolmogorov-Smirnov test). More precisely, there is a higher instance of smaller E-scan amplitudes during movement than during rest (Fig. 29B).

	Rest	Movement	Difference %
λ	1.32	0.74	43.69
D	13.2×10^{-8}	9.41×10^{-8}	28.96

Table 5: Comparison between rest and movement SDE parameters.

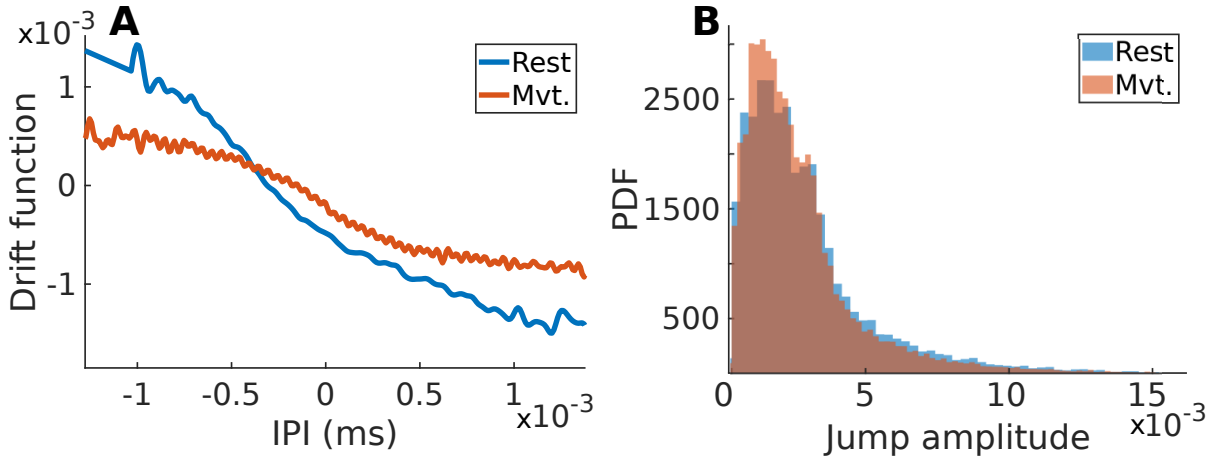


Figure 29: *IPI fluctuations during exploratory movement evolve in a wider potential well than during rest.* **A)** Estimated drift function of the IPI fluctuations during rest (blue) and during exploratory movement (red). Means were subtracted on the x -axis to show both functions on the same graph (mean IPI for rest is ~ 18 ms, and for movement is ~ 15 ms). **B)** Estimated jump amplitude distribution for the E-scans during rest (blue) and during exploratory movement (red).

4.2.2 Oscillatory active sensing

In addition to the stochastic fluctuations analyzed above, the IPI time series also displays bouts of oscillatory activity, referred to as O-states. By inspecting video recordings, we observe that these O-states occur when fish perform repeated back-and-forth movement in front of the few nylon locking nuts located on the periphery of the tank. Although we have access to a few video recordings where fish display O-states, a detailed analysis of the relation between body position and EOD activity is beyond the scope of the current work (see Future work section in the Discussion).

Here we show that the frequency content of O-states consists of two bands centred around 0.5 and 1 Hz, and that the relative spectral power between these bands varies between O-states. Within the same movement bout, we can thus find O-states with different frequency signatures, ranging from (almost) purely 1 Hz to (almost) purely 0.5 Hz, with a tendency for ~ 1 Hz oscillation epochs to have smaller amplitudes than ~ 0.5 Hz ones (Fig. 30A). Contrary to the IPI segments analyzed above, very few E-scans occur during O-states.

We first develop an O-state detection scheme that performs a binary classification between non-oscillatory and oscillatory activity, regardless of the oscillation frequency (see the Methods section for full details). We base this detection scheme on a time-frequency representation of the IPIs during each movement bout. The continuous wavelet transform (CWT) provides a suitable tool for this task, as it gives control over both the frequency and the temporal resolution. This allows for a flexible and pragmatic calibration of the detection procedure. On a time-frequency graph, i.e., the magnitude of the wavelet coefficients, O-states appear as extended bands of elevated magnitude, around either 0.5 or 1 Hz, or a combination of both (Fig. 30B).

Next, for each time, we measure the amplitude of the spectral peaks located within the two specified frequency bands, relative to the surrounding fluctuation ($P_1(t)$ and $P_2(t)$, Fig. 31A). For each peak, this measure, called the peak prominence, is then expressed as a function of time, with the goal of applying a threshold on it to detect moments of oscillatory activity. This approach faces two problems. First, using only information contained within the 1 Hz frequency band makes it difficult to distinguish the small-amplitude 1 Hz oscillations from the surrounding stochastic fluctuations. Second, large E-scans occurring outside O-states yield a strong frequency response (e.g., between 8 and 9 minutes in Fig. 30B) that would inevitably bias the detection procedure if a threshold were simply applied to the peak prominence.

We address both problems simultaneously by scaling the prominence measures as $\rho_i(t) = P_i(t)/\sigma(t)$, $i = 1, 2$, where $\sigma(t)$ is the moving standard deviation (30 s window) of the associated IPI segment. Because 1 Hz oscillations tend to be of smaller amplitude (i.e., small $\sigma(t)$), this allows the weak 1 Hz signal to be enhanced. At the same time, $\sigma(t)$ is sensitive to E-scans, namely, it is increased by E-scans. The net desirable result

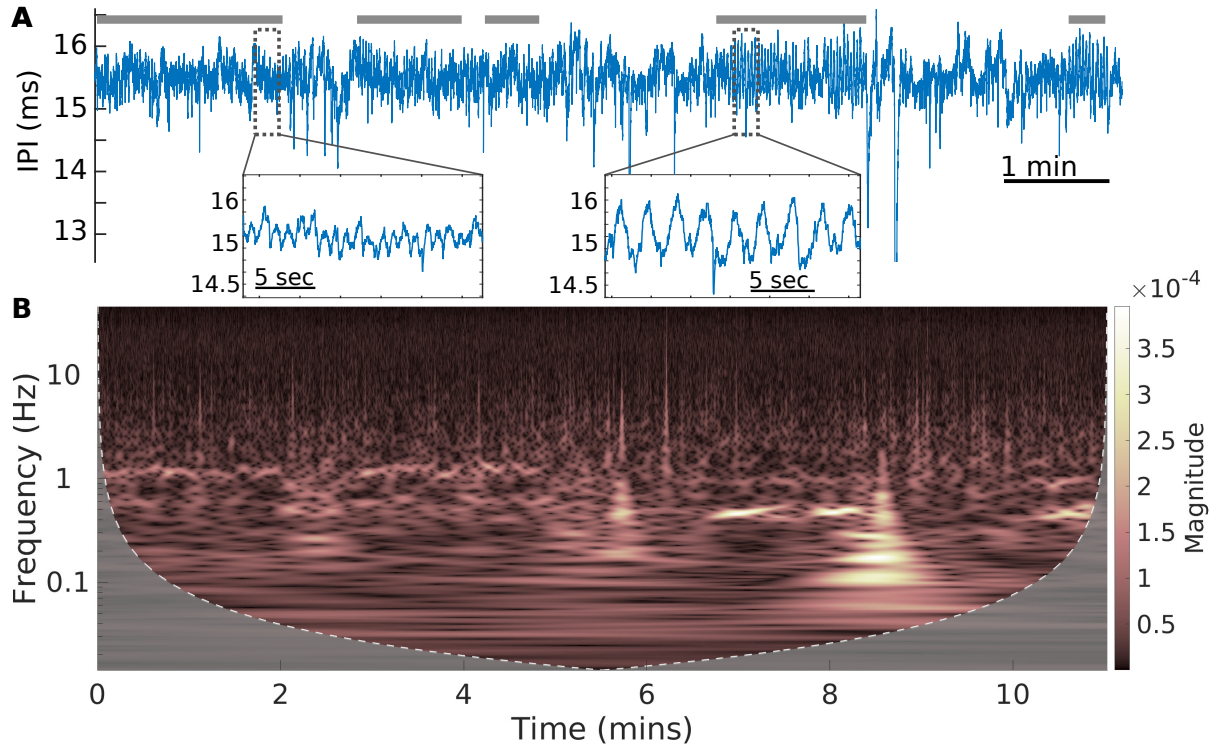


Figure 30: *The relative contribution from two dominant frequency bands varies across O-states. A)* IPI recording showing several O-states within a single movement bout. Grey bars indicate O-state duration, as defined by our detection scheme. Insets show two exemplar O-states with different frequency content. **B)** Time-frequency graph of the same IPI recording showing the magnitude of the continuous wavelet transform. O-states correspond to extended frequency bands with elevated magnitude.

is that this scaling suppresses the contribution of E-scans to the spectrum, enabling the presence of the 0.5 and 1 Hz rhythms to be highlighted. After choosing the appropriate threshold (see Sec. 4.4.3 for details), we detect an O-state when either or both $\rho_1(t)$ and $\rho_2(t)$ are above threshold. (Fig. 31). This scheme yields an accurate binary classification between oscillatory and non-oscillatory activity of the IPI.

By averaging $\rho_1(t)$ and $\rho_2(t)$ over the duration of each O-state, we can quantify each O-state by a pair, $(\bar{\rho}_1, \bar{\rho}_2)$, and thus visualize the oscillatory activity of the entire IPI time series on a 2D scatter plot (Fig. 32A). The appearance of two clusters in the resulting plot suggests the existence of two preferred modes of oscillatory sampling. A typical O-state in cluster 1 shows a dominant frequency at 1 Hz with smaller component at 0.5 Hz, and vice-versa for the O-states of cluster 2 (Fig. 32B). In addition, the movement variable during most O-states shows a dominant 0.5 Hz frequency. There is, however, a

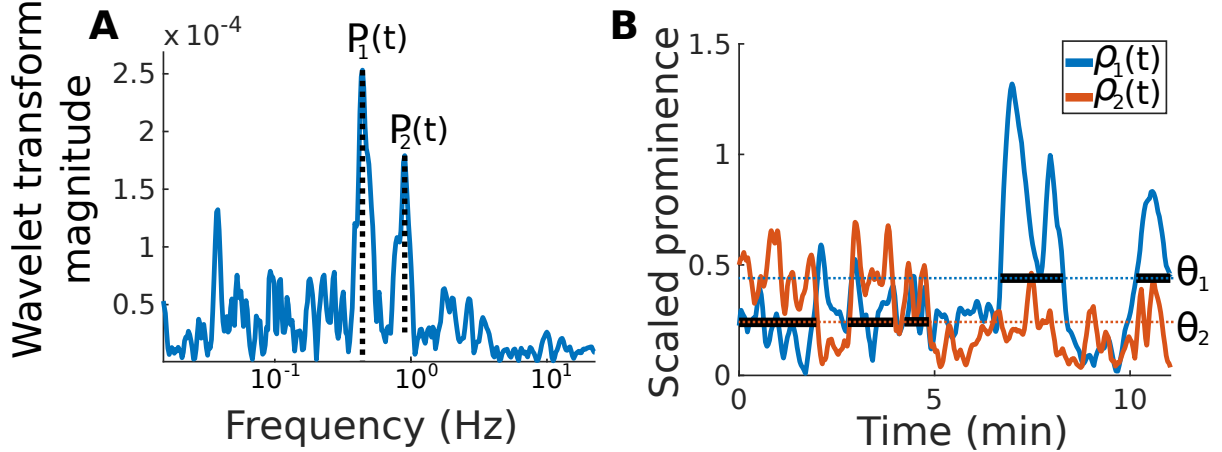


Figure 31: *O*-state are detected by applying a threshold to the scaled prominence measures. **A)** A slice of the wavelet spectrum of Fig. 30B at 10.5 min. Peaks are clearly seen around 0.5 Hz and 1 Hz. $P_1(t)$ and $P_2(t)$ (vertical dotted line) quantify the prominence of those peaks, respectively at 0.5 Hz and 1 Hz, relative to the surrounding fluctuations. **B)** After dividing the prominence of each peak by the moving standard deviation, we obtain two new measures, $\rho_1(t)$ and $\rho_2(t)$, that reliably capture the appearance of oscillations in both frequency bands. We apply the thresholds θ_1 and θ_2 on $\rho_1(t)$ and $\rho_2(t)$, respectively.

group of *O*-states in cluster 2 where the dominant movement is at 1 Hz instead (black dots, Fig. 32).

O-states in cluster 2 typically show more coherent oscillations than those in cluster 1. This can be seen, for instance, from the autocorrelation function of exemplar *O*-states from both clusters (Fig. 33A). To quantify this difference, we calculate the so-called characteristic time [155]:

$$\tau_c = \int_0^\infty C^2(t) dt \quad (86)$$

where $C(t)$ is the (normalized) autocorrelation function of the IPI times series (with their means values subtracted). Larger values of τ_c are associated with more coherent oscillations. Across all *O*-states, we find a positive correlation between the characteristic time and $\bar{\rho}_1$ (Fig. 33B), but no significant correlation with $\bar{\rho}_2$.

Individual *O*-states can also be visualized as a density in phase space, where trajectories are depicted as a function of both the IPI and its derivative. We see, for instance, that trajectories of *O*-states from cluster 1 (Fig. 34A) are distributed around the origin more symmetrically than those of cluster 2 (Fig. 34B). This indicates that the IPIs




Figure 32: *O*-state appear as two distinct clusters in $\bar{\rho}_1$ - $\bar{\rho}_2$ space. **A)** Each *O*-state can be quantified by its mean scaled prominence measures, and thus represented by the pair $(\bar{\rho}_1, \bar{\rho}_2)$. *O*-states whose associated movement variable shows a dominant frequency around 1 Hz (0.5 Hz) are shown as black (red) dots. **B)** PSD of exemplar *O*-states from cluster 1 and 2.

during the slower (0.5 Hz) *O*-states tend to spend more time in the trough than in the crest of the cycle.

4.3 Discussion

We have shown that pulse-type electric fish exploit different dynamical regimes for electrosensory acquisition depending on the behavioural context. During rest and exploratory movement, the IPI time series displays stochastic fluctuations along with the occurrence of E-scans. A dynamical system model of these fluctuations consisting of a nonlinear drift, an additive Gaussian white noise and a jump process could be fitted well to the data. This analysis reveals that the more stable fluctuations during rest can be explained by a potential well that is narrower than during exploratory movement (the drift function is the negative derivative of this potential). By interpreting this result in the context of the phenomenological model of active sensing presented in Ref. [19], we hypothesize that this quantitative difference in the shape of the potential well arises from different levels of reafferent feedback strength, which is thought to change the correlation time (represented by the drift function here) of ongoing neuronal activity.

In addition, we find that the noise intensity and E-scan rate are both higher during rest than during exploratory movements. This could be to compensate for the narrower

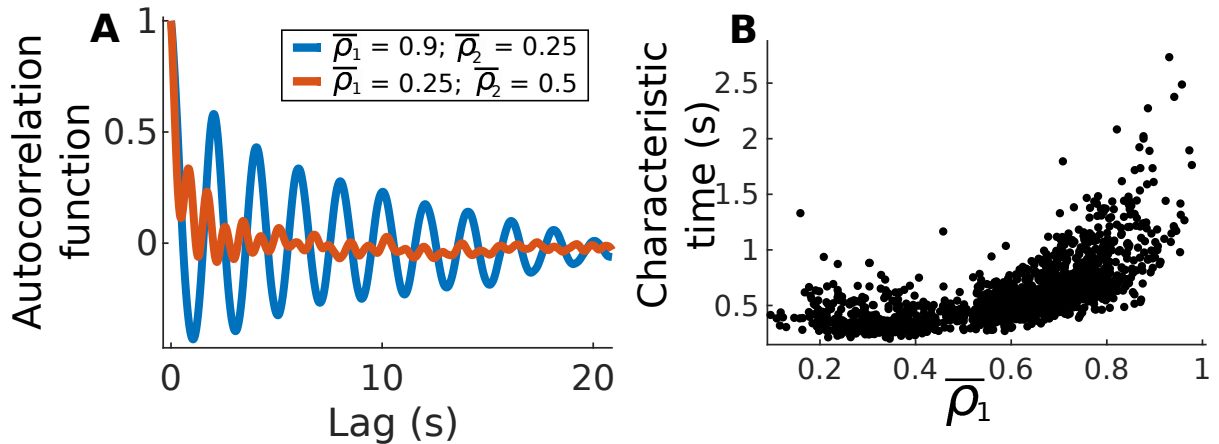


Figure 33: *O*-states from cluster 2 have more long-lived correlations than those of cluster 1. **A)** Autocorrelation function of *O*-states from cluster 1 and 2. Fig. 32A. **B)** The characteristic times (integral of the squared autocorrelation function, Eq. (86)) of *O*-states correlate with $\bar{\rho}_1$ (but not with $\bar{\rho}_2$).

potential well, which suppresses some of the low-frequency content, assuming of course that there is a functional benefit to the jump-diffusion-type fluctuations as opposed to more steady IPIs. Furthermore, when considering E-scan amplitudes, smaller E-scans are only slightly more probable during movements than during rest. This suggests that, despite their higher rate during rest, the core neuronal mechanism underlying E-scans remains largely unchanged.

On the other hand, we also show how fish drastically change their sensory sampling strategies when attending to specific stimuli. Indeed, during *O*-states, IPIs evolve periodically with a variable degree of coherence (as measured by the characteristic time, Fig. 33B), and a variable frequency content (as quantified by the averaged prominence measures $\bar{\rho}_1$ and $\bar{\rho}_2$, averaged over the duration of each *O*-state, Fig. 32A). To our knowledge, this specific sensorimotor pattern has not yet been documented, and should be added to the behavioural repertoire of *Gymnotus sp.*

The two frequency bands (0.5 Hz and 1 Hz) observed in the IPI are also present in the movement variable, suggestive of an underlying coupling mechanism between sensory sampling and position. For instance, corollary discharges associated with motor output could potentially reach the pacemaker nucleus and modulate its activity in a manner that correlates significantly with movement. Another possibility could be the involvement of rhythmic input streams from other modalities, such as the lateral line (which




Figure 34: *O*-states show an ovoid phase space structure with no apparent limit cycle. Phase portrait density of the same 2 *O*-states analyzed in Fig. 33A, i.e., from cluster 1 (**A**), and cluster 2 (**B**).

measures pressure gradients) and proprioception (which reports on body positioning, i.e., tail bending). Due to the back and forth motion, hydrodynamic input to the lateral line is likely to show the same periodicity as the movement itself. Similarly, the specific movement pattern during *O*-state involves a sequence of tail-bending to reposition the body during the cycle. Since proprioceptors reliably code for tail-bend angle [10], they should also displays some version of this periodic activity. It is possible that these rhythmic signals are re-routed to electrogenic centers and impart their frequency content to the IPIs.

It remains to be determined why these fish perform rhythmic movements during *O*-states, since they could easily sense the same stimulus by remaining at a fixed position. There must be a considerable benefit to warrant the energetic cost of such movements. Furthermore, why and how electrogenesis becomes coupled to motor output is an open question. Similar questions have been asked with respect to sniffing and whisking in rodents, both of which are rhythmic and also coupled [38]. We hypothesize that *O*-state are an attentive behaviour during which fish perform movements to enhance their perception of the attended stimulus. Rhythmic movement can indeed provide better sensory representations of input acquired at specific phases of the cycle [172]. In rodents, the motor rhythm associated with sniffing and whisking influence sensory encoding in both the barrel cortex [32] and the olfactory bulb [176]. In the primate visual system, neuronal excitability is increased at the onset of eye fixation [161], and the alternation between fixations and abrupt saccades occur approximately at a 3 Hz rhythm [172]. Overt motor

rhythms can enhance sensory processing even when the sensor is not directly affected by the movement, such as for auditory processing [138, 139, 137]. There may also be a prediction-verification involved during O-states, i.e. a predictive coding whereby moving closer and further away from an object while sampling it at waxing and waning rates likely produces expected sensory signals, and deviations from these expected signals caused by novel, unpredictable stimuli may be highlighted.

It has been suggested that when no rhythmic stimuli are present, attention operates in a continuous mode, associated with a vigilance state [171]. During such a state, low-frequency neuronal oscillations are suppressed, as they would be detrimental for the detection of unpredictable stimuli. In contrast, rhythmic stimuli enable attentional processes to enforce entrainment of low-frequency oscillation in the delta band [109, 172, 22]. By engaging in low-frequency motor rhythms, such as the O-states reported here, animals might be able to leverage this attentional process. In addition, recent evidence in humans and macaques suggests that even sustained attention is in fact modulated by ongoing, low-frequency neuronal excitability fluctuations [74, 50]. In our case, perhaps an underlying attentional rhythm is always present, but is only overtly manifested during O-states. Active sensing is indeed an overt expression of what an animal is attending to [197, 38, 135, 97]. It is therefore possible that the rhythmic sensorimotor strategy used by fish during O-states is a behavioural readout of an equally rhythmic attentional process.

4.3.1 Future work

Our result establishes a hitherto unknown two-way link between different features of active sensing, in the form of an oscillatory sensory sampling rate and rhythmic movement. Our analysis of this phenomena opens many questions. With further analyses and experiments, O-states could perhaps serve as an exciting new probe into the link between active sensing, motor rhythms, and attention. In terms of data analysis, the next step would be to determine the precise coupling between movement and the IPIs. With image processing tools, the head position and body orientation can be extracted from video recordings. What is the phase relation between the IPI and movement parameters? Is it dependent on the frequency band? Is there phase-amplitude coupling between those frequencies? Applying bandpass filters centred on the frequency bands of interest here and applying the Hilbert transform would enable us to obtain this information. Another

potentially fruitful avenue is whether the transition in and out of O-states is gradual or abrupt. By closely analyzing the data around the transitions, we could determine if the electrosensory and the motor rhythms emerge already in phase, or if one gradually entrains the other.

Additional experiments would also yield valuable insight into this phenomenon. For instance, it has been shown that fish exploit rhythmic back-and-forth movements to enhance their sensory representations during refuge-tracking tasks [183], and that this behaviour is under closed-loop control [17]. Although these studies were performed on wave-type electric fish, and were specifically designed to monitor longitudinal movements, the observed behaviour might have a similar function and origin as the O-states reported here. Unfortunately, only motor output were recorded in these studies, so any causal link to the EOD dynamics for these fish will have to be investigated with new experiments.

Here we show that rhythmic sensory sampling, in addition to the already rhythmic motor routine, adds another layer of complexity to the spatiotemporal sensory flow during O-states. This begs the question of whether the IPIs are under a similar, or the same, closed control loop as during the refuge-tracking tasks. By implementing an online video tracking system, one could detect when fish perform O-states and dynamically move the attended stimulus in and out of phase with the movement pattern. If fish modulate the amplitude of their movement as a function of the phase difference between their movement and that of the stimuli, this would suggest that reafferent feedback is indeed used by a controller to regulate sensory inflow. Crucially, monitoring what happens to the phase and amplitude of the IPI could shed some light on the nature of their coupling with the motor output. Furthermore, if attention is indeed rhythmically allocated during O-states, we would expect that perceptual sensitivity to a distracting signal, such as a sound cue, would vary as a function of the phase of the observed O-state cycle.

4.4 Methods

4.4.1 Data preparation

The raw data consist of EOD time-stamps. A sequence of IPIs is obtained by taking the difference between successive EODs. The movement variable is calculated as per

the method developed in Ref. [96]. Both the IPI and the movement variable are then resampled as 100 Hz. Experiments for the particular fish analyzed here are divided into 11 recording sessions, for a total of approximately 110 hours of recordings. All sessions were concatenated, and a histogram for the overall movement variable is obtained. From this histogram, we observe that the movement variable is unambiguously bimodal. A threshold is thus applied to determine whether the fish is in a moving or a resting state.

For all analyses carried out here, we choose only states that are longer than 100 s. This is to allow for a sufficient amount of data away from the transitions between moving and resting states, which tend to show non-stationary characteristics [132]. To avoid these transients, we also discard 5 s on both sides of those transitions. In addition, the IPI during certain moving or resting states show a clear long-term trend, such as a gradual stabilization after the offset of movements. We remove these trends by subtracting the moving median (50 s window), which is less sensitive to outlier values caused by E-scans than the moving average. This process yields a set of long, detrended IPI time series corresponding to various resting or moving bouts.

4.4.2 Jump-diffusion modelling

After concatenating all resting state IPIs, as prepared by the method outlined in Sec. 4.4.1, we apply the inference scheme developed in Ref. [131] to estimate the parameters and function of Eq. (84). A threshold for jump (E-scans) detection is first established based on the statistics of the increments of the IPI time series. Jumps are then detected and used to obtain an initial estimate of the rate, λ , and amplitude distribution, Q_A . In addition, by applying a range of threshold values on the original data, and by analyzing the diffusive segments between detected jumps, an estimate of the noise intensity, D , is produced. These estimates are used in the differential Chapman-Kolmogorov equation [57] to yield a first estimate of the drift function, F . Simply applying a threshold to a jump-diffusion process, however, introduces systematic errors due to the presence of false positives, i.e., large diffusive fluctuations erroneously identified as jumps. Consequently, an iterative scheme is applied where these false positives are gradually demixed from the initial pool of detected jumps, leading to a reduction of the bias in the initial estimates.

4.4.3 O-state detection

Identifying O-states first relies on the creation of two time-dependent variables that vary according to the frequency content found in the IPI time series. We then detect O-states by applying a threshold on those variables. To achieve this goal, we proceed by obtaining a time-frequency representation of the IPIs for each movement bout separately. We choose the continuous wavelet transform, with Matlab’s bump wavelet at 32 voices per octave (function “cwt” in the Wavelet toolbox). This particular choice of wavelet yields a better frequency resolution, but poorer temporal resolution, compared to the more common Morlet wavelet. This combination is desirable since the two frequency bands of interest (0.5 and 1 Hz) are relatively close to each other. From visual inspection of the overall power spectrum of the IPI during movement, we define these two bands by the intervals [0.3, 0.65] Hz, and [0.8, 1.4] Hz. Defining the two frequency bands in this way gives the necessary flexibility in the detection scheme, as different O-states have slightly different frequency content.

Next, for each time within the current movement bout, we find the location of the maximum of the CWT magnitude in both frequency bands and calculate the associated peak prominences (Fig. 31A). If there are no peaks within the intervals (such as during most stochastic fluctuations), we set the prominence to 0. Both prominence measures are then divided by the moving standard deviation (30 s window) of the current IPI trace, the result of which is smoothed with a 10 s moving average window. This process is applied to each movement bout, and all the resulting pairs of scaled prominences are concatenated in order to obtain a global histogram. These two histograms are unimodal, but show a clear maximum in their second derivative (Fig. 35) which can be used as appropriate thresholds ($\theta_1 = 0.45$, $\theta_2 = 0.25$).

Lastly, we apply these thresholds to the scaled prominence in order to obtain two binary variables that represent whether a frequency band is active or not at any given time. We then define O-states as moments when either ρ_1 or ρ_2 , or both ρ_1 and ρ_2 are above threshold. O-states shorter than 25 s are discarded to ensure that we do not include mere fluctuations in our analysis.




Figure 35: Histograms of the scaled prominence are used to define the *O*-state detection thresholds. The scaled prominence for both frequency bands, $\rho_1(t)$ (0.5 Hz) and $\rho_2(t)$ (1 Hz) are concatenated across movement bouts and compiled in the histograms shown here. The thresholds are defined as the value of the scaled prominence for which the second derivative of the histograms are at a maximum.

Modelling the electrosensory signature of endogenous vs. exogenous chirps in electric fish

5.1 Introduction

Social encounters among wave-type electric fish produce a seemingly confusing sensory scene [189]. Indeed, when several fish are in the vicinity of each other, the electric fields produced by their electric organ discharge (EOD) summate. This cumulative effect can be understood as a sum of dipole-like sources, each oscillating at the frequency of the corresponding EOD. If two fish have EOD frequencies that are too similar, they automatically raise or lower their own frequencies such as to increase the frequency difference. This is a well-known behaviour called the *jamming avoidance response* [73]. The specifics of this response vary across species.

The summation of dipole-like sources creates a rich spatial distribution of the electric potential. The oscillating nature of these sources, however, also creates a temporal

beat pattern that slowly modulates the amplitude of the faster electric potential oscillations, hence, we refer to this beat as the amplitude modulation (AM). The AM itself is described by its frequency, $\Delta f = |f_1 - f_2|$, and its amplitude, which varies in space depending on the relative location of each fish. Importantly, electric fish sense the oscillating electric potential along their body¹¹ and are able to extract the AM, which contains behaviourally relevant information [133]. Any changes to the AM frequency or amplitude are detected by the fish and, as such, can be used as a means of communication.

In social contexts, such as courtship or aggression, electric fish emit communication signals in the form of frequency and amplitude modulation of their own EOD [87, 88, 75]. Many different types of signals exist, but here we focus on the small chirp (also known as Type 2 chirp), a transient excursion of a fish’s EOD frequency. These events last around 15-20 ms, during which time the frequency increases by 50-100 Hz. Because they briefly affect the frequency difference between each fish’s EOD, Δf , chirps also have an impact on the AM pattern, specifically on its phase. Indeed, the AM displays a phase advance (also called phase reset) during a chirp, the magnitude of which depends on the frequency difference at the peak of the chirp and on the AM phase at chirp onset [210, 13, 134]. This phase advance is detected along the body of the fish and is believed to quantify the sensory information relevant to the chirp signal.

Curiously, when fish are engaged in communication through chirps, it is unclear how they recognize their own chirps (endogenous) from those of other fish (exogenous), as both cases can, a priori, exert similar changes on the beat phase [175]. Importantly, electric fish belonging to the *Gymnotiform* order, such as those studied in this thesis, are thought not to possess a neural copy of the outgoing signals sent to the electric organ to elicit EODs. This copied signal, better known as the *corollary discharge*, is a strategy found across the animal kingdom to differentiate sensory signals that result from one’s own actions (reafference), such as movements or EODs, and those that originate externally (exafference). If indeed there is no intrinsic neural mechanism for endogenous vs. exogenous chirp differentiation, then perhaps this is achieved through some sort of sensory signature.

¹¹More precisely, they sense the transdermal potential difference at specialized sites on their skin called electroreceptors.

5.2 EOD model

This line of questioning motivates the modelling presented below, where we apply a previously established EOD model [5, 103] to examine the spatial distribution of chirp-induced AM phase resetting along the skin of the wave-type fish *Apteronotus leptorhynchus*. In this model, the spatial component of the EOD originates from a data-calibrated pattern of current source density along the electric organ (Fig. 36A). This was previously implemented in the finite-element modelling software COMSOL Multiphysics as a two-dimensional, electrostatic boundary-value problem [5] (see Appendix A). By solving the Poisson equation for a given phase of the oscillation, the result of this specific EOD is a dipolar-like electric potential function consistent with experimental observations [103]. When two fish are introduced in the simulation domain, we can observe the cumulative electric potential for any given phase (Fig. 36B).

Temporally, the EOD oscillates at ~ 700 Hz with some harmonics (Fig. 36C), and the presence of a second fish with a different frequency creates a rich beat pattern when the two EODs interfere (Fig. 36D). The amplitude of this beat pattern varies according to the relative distance and the orientation of two nearby fish. In what follows we neglect the harmonics and assume the EODs are pure sine waves. This facilitates the extraction and interpretation of the AM phase. Future work will evaluate in greater depth the impact of these harmonics on the chirp signature.

5.3 Chirp model

Following Ref. [13], chirps are modelled as Gaussian frequency excursions, $f_{chirp}(t)$, where we use a peak frequency difference, f_{peak} , of 80 Hz (Fig. 37A). The variance of the Gaussian is set by the duration of the chirp, $\Delta t=15$ ms. This frequency excursion is added to the baseline (constant) frequency, f_1 . The phase of the EOD of the chirping fish is calculated as

$$\varphi_1(t) = 2\pi \left(f_1 t + \int_0^t f_{chirp}(t) dt \right),$$

and the EOD waveform varies as $\cos(\varphi_1(t))$ (Fig. 37B). When this is added to another EOD waveform with a different (constant) frequency, the resulting AM shows the aforementioned phase advance during the chirp (Fig. 37C).

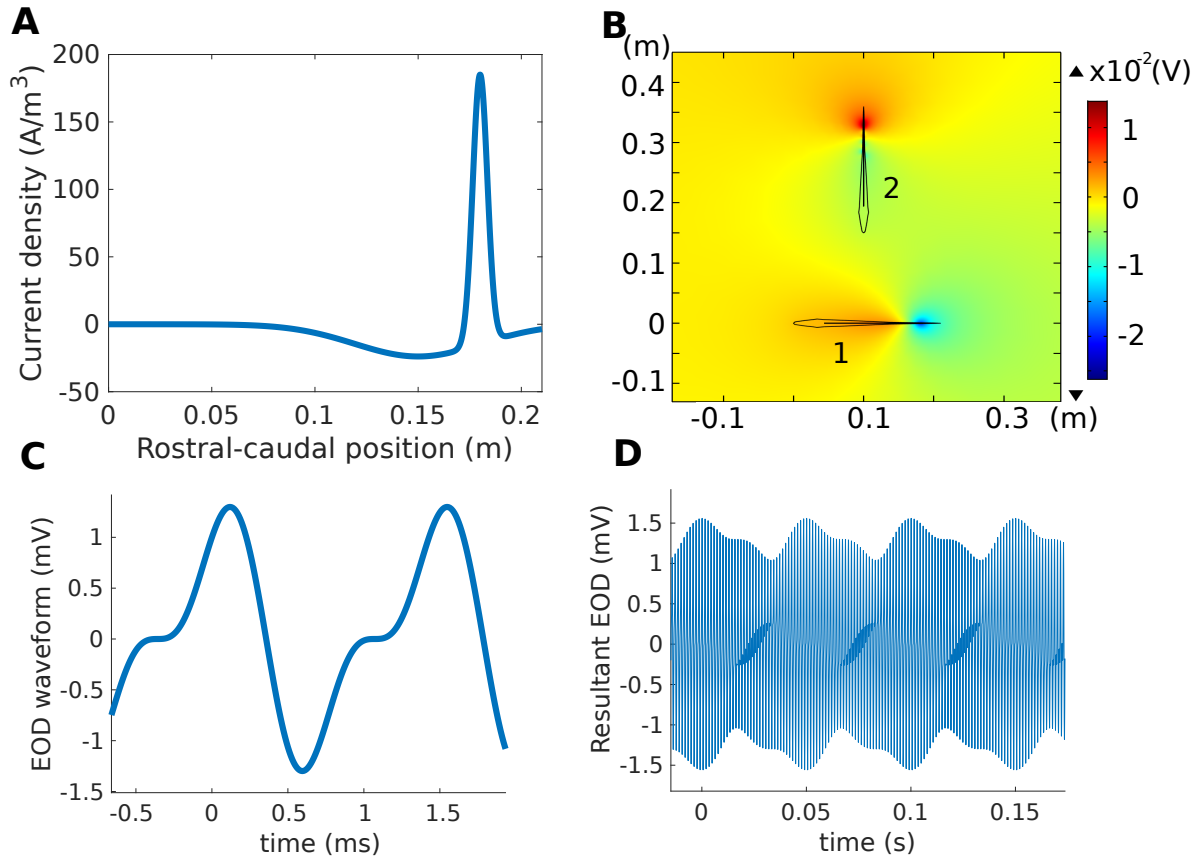


Figure 36: **A)** Current density along the electric organ during the tail-positive phase of the EOD (head = 0 m; tail = 0.2 m). The estimate of this density profile was inferred from a computational analysis that used both COMSOL Multiphysics and a genetic algorithm [103]. **B)** Snapshot of the electric potential, with respect to a reference electrode located at the edge of the tank (outside the depicted frame), resulting from the interaction of two fish with their individual EODs. **C)** Model of the EOD waveform of *Apteronotus leptorhynchus*: $\cos(2\pi f_1 t) + 0.5 * \cos(4\pi f_1 t - \pi/2)$, with $f_1 = 700$ Hz. **D)** Beat pattern resulting from the superposition of two EOD waveforms with $\Delta f = 20$ Hz.

To quantify this reset, we need to extract the instantaneous phase of the AM itself, i.e., the phase of the beat pattern as a function of time. The common technique to achieve this is by constructing the so-called analytic signal from the Hilbert transform¹². A first transform is applied directly to the resultant EOD signal, i.e., the “carrier wave” (blue curve, Fig. 37C), and the amplitude of the analytic signal gives the AM signal, or “envelope” (red curve, Fig. 37C). A second Hilbert transform is now applied to the AM itself and used to construct a second analytic signal, the argument (i.e., angle) of which gives the desired AM phase. The presence of the chirp, and its influence on the

¹²The analytic signal is a complex function built from a real signal. Its real part is simply the original signal, while its imaginary part is the Hilbert transform of the original signal.

AM makes it impossible to read off the AM phase directly from the AM signal, hence the need for the second transform. The extracted AM phase clearly shows how the phase advances more rapidly during the chirp, which can be quantified by $\Delta\varphi_C$, the difference between the AM phase before and after the chirp (Fig. 37D).

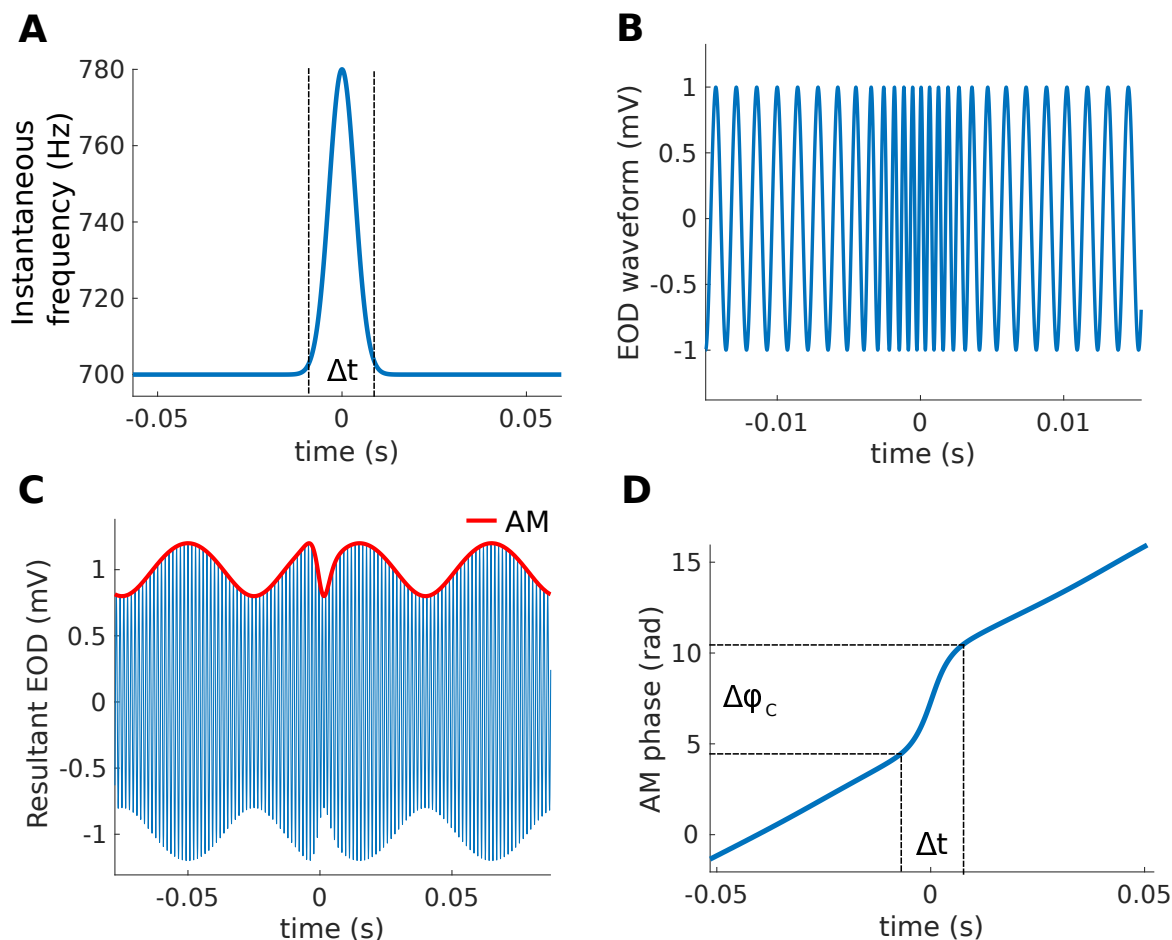


Figure 37: **A)** Gaussian frequency excursion used as a model for chirps, $f_{chirp}(t)$. **B)** Sinusoidal EOD waveform during a chirp. For the sake of clarity, the trace shown here is the result of an exaggerated version of the excursion in A. **C)** Superposition of two EODs with $\Delta f=20$ Hz. One of them chirps at $t=0$ s, which results in a phase advance of the AM (red curve). **D)** Instantaneous phase (unwrapped) of the AM in C, extracted from two applications of the Hilbert transform to the resultant EOD in C. It shows the phase advance $\Delta\varphi_C$ during the chirp of duration Δt .

5.4 Results

At any given time, the AM phase varies along the length of a fish due to spatial variations of the phase of its own EOD (e.g., the head and tail region are fully out of phase with each other). For this reason, the impact of a chirp on the AM will also vary along the length of the fish (Fig. 38A). This is quantified in *phase reset curves*, where $\Delta\varphi_C$ is plotted as a function of the position on the body of the fish (Fig. 38B). This analysis then culminates in the question of whether the information contained in this phase reset curve depends on which fish chirped, i.e., endogenous vs. exogenous chirping. More precisely, *does there exist a sensory ambiguity whereby different chirping scenarios elicit the same phase reset curve?* If so, then fish would not be able to reliably depend on the sensory signature of chirps to discriminate between an endogenous and exogenous origin.

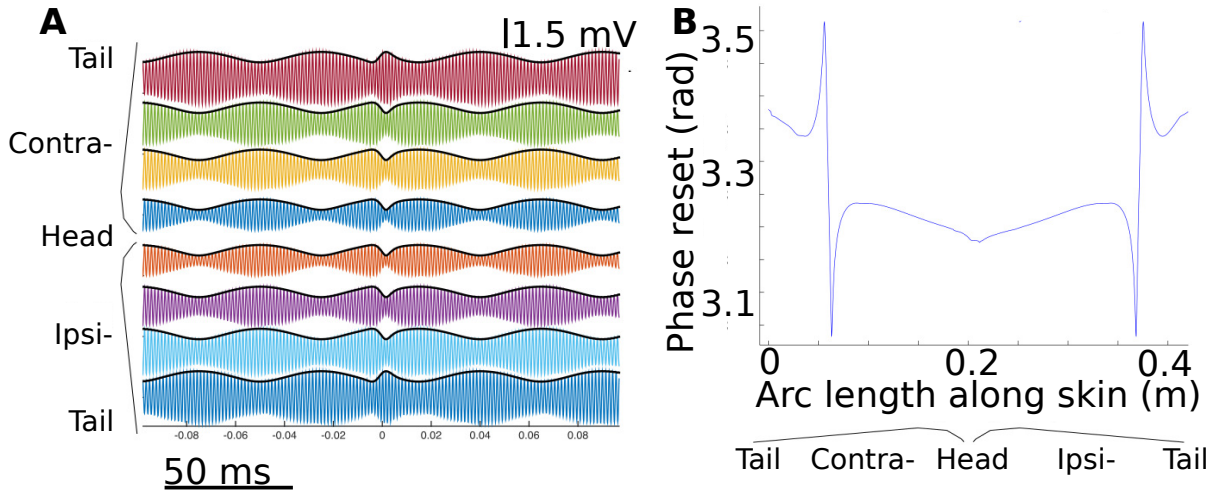


Figure 38: **A)** Beat patterns at different locations along the length of the fish (fish 1 in Fig. 36B), shown for both the ipsilateral and the contralateral side to the other fish (fish 2). **B)** Phase reset curve showing $\Delta\varphi_C$ as a function of the location along the body of the fish who chirped (fish 1 in Fig. 36B). The configuration used in this analysis, i.e. fish separated by ~ 15 cm, perpendicular to one another (Fig. 36) and with $\Delta f = 20$ Hz, yields very little asymmetry between both sides.

To evaluate this ambiguity, we run two sets of simulations, one corresponding to endogenous chirps and the other to exogenous ones. In both sets, we run an ensemble of simulations of the situation described above (i.e., two fish, separated by 15 cm, perpendicular to each other, one of which chirps with a peak frequency $f_{peak} = 80$ Hz for 15 ms), where a different Δf is used for each simulation. For the endogenous set, fish 1 chirps, and for the exogenous set, fish 2 chirps. The AM is measured along the body of fish 1

for both sets.

We calculate the phase reset curve for each simulation of the ensembles. This results in two sets of phase reset curves, both indexed by values of Δf , that can then be compared with a given metric (here we choose the root-mean-square error). This comparison can be summarized on a two-dimensional grid, where the values represent the “distance” between two given curves (Fig. 39). This grid thus quantifies how the AM phase reset curve along the body of the fish for an endogenous chirp differs from that for an exogenous one. From this figure, we observe cases where the phase reset curves are indistinguishable, or at least very similar, indicative of sensory ambiguity. Indeed, for an endogenous chirp occurring on a beat with, e.g., $\Delta f=20$ Hz, there exists an equivalent exogenous chirp, occurring on a beat with $\Delta f=-20$ Hz (and other higher multiples) that produces the same sensory signature. These chirp signatures are therefore degenerate, in the sense that two (or perhaps more) scenarios yield the same sensory information.

To be clear, if we consider a single encounter between two fish, where Δf is fixed at, say, 20 Hz, the sensory signature of an endogenous chirp will be significantly different from that of an exogenous one (yellow/green areas of Fig. 39). This is because the instantaneous frequency difference at the peak of the chirp is not the same in both cases: when the higher frequency fish chirps, the maximum frequency difference during the chirp is $f_{peak} + \Delta f$, while it is $f_{peak} - \Delta f$ if the lower frequency fish chirps. In other words, a specific encounter between two fish does not lead to any sensory ambiguities. Rather, the ambiguity stems from the fact that there exists other possible encounters that share the same sensory signature. This results motivates the search for a neural correlate of chirp production in the ELL.

5.5 Future work

In summary, our modelling framework provides new supporting evidence for the presence of a corollary discharge associated with the production chirps in *Apteronotus leptorhynchus*. This idea is fully examined in a larger manuscript (in preparation) where the current analysis will appear. In addition to the current analysis, further work would be necessary to better characterize how phase reset curves change across various conditions. For instance, it would be appropriate to verify how the shape of Fig. 39 changes with

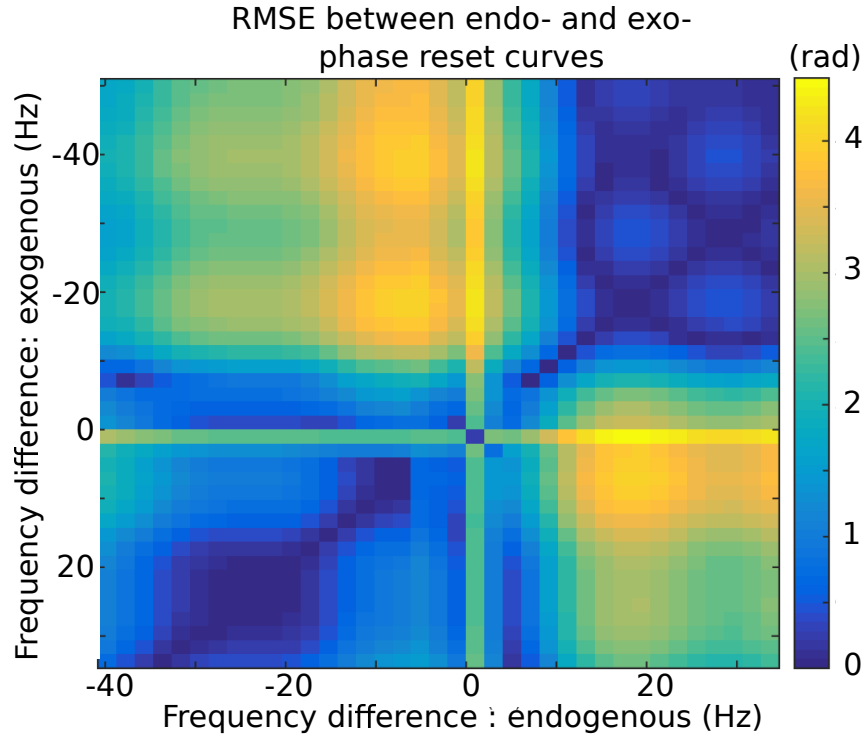


Figure 39: Comparison of phase reset curves between endogenous and exogenous chirp scenarios. The distance between two curves is measured as the root-mean-square error. Dark blue regions correspond to scenarios where there is little to no difference in phase reset curve between endo- and exo-chirps, leading to a sensory ambiguity as to which fish produced the chirp.

respect to changes in the parameters of the simulation (e.g., f_{peak} , Δt , distance between fish, angle between fish). Furthermore, here we neglect the harmonics present in the individual waveforms, but it would be necessary to determine the impact of including them. To partially address these questions, it would be enlightening to re-create the same analysis presented here, but with simple dipole sources rather than the COMSOL implementation of the detailed EOD waveform, which requires a computationally expensive finite element solver. Such a toy model would be computationally much lighter and allow for a more efficient exploration of the questions at hand. Perhaps this toy model could further serve as a means to examine chirping scenarios where more than two fish are present, or to examine whether the concept of phase-reset curves can be used to quantify the information content of small chirps. This also raises the related question of whether information transfer is maximized in a certain region of the parameter space.

5.6 Appendix A: COMSOL implementation

The simulations are performed in the finite element software COMSOL Multiphysics. The simulation domain consists of a 2 m by 2 m region of space meshed with the “Very fine” setting. The electric fish model introduced in this domain is morphologically accurate and consists of a thin rectangular electric organ where the current density is implemented (Fig. 36A), an interior compartment, and a skin layer, all with their respective conductivities. With both fish implemented in the model domain, we run electrostatic simulations to obtain the electric potential function (with respect to a reference electrode located on the periphery of the domain) at a given phase of the oscillation. A parameter sweep is applied to index the time parameter, which creates a series of solutions for all the phases of the oscillation, and over several cycles. The time series for the electric potential along various points along the skin are then extracted for post-processing (i.e., extracting the AM). Finally, a second parameter sweep is performed over the values of Δf in order to generate the endogenous and exogenous sets mentioned earlier.

Conclusion

In this thesis, I explore how low-dimensional stochastic models can provide meaningful representations of the spontaneous activity of complex systems. I apply this idea at two vastly different scales of neural activity, namely at the single-neuron and at the behavioural level, and I demonstrate the value of this approach in advancing our intuitive understanding of these systems.

Summary and outlook

In Chapter 2, I propose that the long-term movement statistics of freely behaving fish can be explained by the interaction of a slow forcing with a noisy bistable system. The slow external forcing modulates the transition rate between stable states and produces stretched exponential residence time distributions that fit well with those of the experimental data. Furthermore, I exploit the strong coupling between the EODR and the movement variable to analyze finer details observed around transitions. Notably, the EODR time-course during transitions toward the resting state, following a long active state, is consistent with a slow tilting of a stochastic double-well dynamical system. This process involves two saddle-node bifurcations, where the system is taken from the initial monostable active state, through a transient bistable period, and then arrives at the monostable resting state. Similarly, the variability observed in the shape of the transitions from resting to moving can be explained by a coupling between the slow external forcing and the parameters of the potential function.

Considerable insight can be gained when animals freely explore their behavioural repertoire [14]. My analysis, for instance, hints at a unified framework to understand intermittent locomotion, and possibly behavioural state transitions in general, across many types of organisms. With advances in high-throughput monitoring of spontaneous animal behaviour [14, 41, 148, 77, 7], there is indeed a need for such interpretative frameworks. Due to its generality, the effective model I propose could encompass dif-

ferent neuronal implementations of the same overall dynamics (e.g., different network properties, neuromodulators, etc.). This is exemplified in the fact that, with changes in only a few parameters, the proposed model fits the data from two fish with vastly different behaviours. Such detailed analyses of spontaneous animal behaviour is required to fully understand the impact of internal states on responses evoked by external stimuli [164]. Effective models can thus act as tools to assist our understanding of how internal states shape behavioural output, and perhaps lead to a better characterization of the impact of pharmacological and neuromodulatory agents on behaviour.

In chapter 3, I develop a data-driven inference method to estimate the parameters and functions of a jump-diffusion stochastic process. Central to my approach are the probabilistic calculations associated with false positives: diffusive fluctuations erroneously detected as jumps. These false positives unavoidably appear when applying a threshold to the increments of a sampled jump-diffusion process. I show that the probability of observing false positives, as well as their amplitude distribution, are themselves stochastic functions, as they depend on the current value of the process. By accounting for this in the analysis, I obtain accurate corrections for the observed (and biased) jump rate and amplitude distribution.

I then apply this method to recordings of spontaneous fluctuations observed in the membrane voltage of ELL pyramidal cells. I reconstruct effective jump-diffusion models corresponding to several holding currents, which provides a dynamical explanation for how different types of pyramidal cells change their voltage fluctuation statistics in response to depolarization. Notably, I found that, as a function of depolarization, CLS cells increase their blip rate, but not blip amplitude, and that CMS cells, despite having no apparent blips, show an increase in diffusive noise intensity.

Intrinsic membrane voltage fluctuations are the substrate on which synaptic input are integrated [206]. With a proper dynamical representation of these fluctuations as a function of cell type and depolarization level, the functional role of blips in electrosensory processing could be investigated. The successful fit between the proposed jump-diffusion models and the experimental data allows for this possibility, despite the unknown biophysical origin of blips. Indeed, given the current body of knowledge on the ELL circuitry [16] and on the neural basis of electrosensory encoding [144], one can envision a model

where realistic, context-specific sensory input is fed to an array of pyramidal cells whose intrinsic activity is controllable through the parameters of the effective models. Understanding the relation between network output and the statistics of intrinsic fluctuations could potentially unveil a network-level computational role of membrane-bound biophysical processes.

In Chapter 4, I return to the EODR time series of freely behaving fish, although here I use the equivalent representation in terms of IPIs. I apply the same jump-diffusion modelling approach as in Chapter 3 to the IPI fluctuations on a shorter timescale than in Chapter 2. These fluctuations display randomly occurring E-scans and are well-fitted by effective jump-diffusion models. I find that the statistical difference between the IPIs time series during rest and exploratory movements can be mostly explained by a narrower and less symmetric potential function during rest than during movement. Given the known implication of E-scans in food-searching tasks, a proper estimation of the E-scan rate and amplitude statistics is desired, as this can help evaluate potential neural mechanisms involved in their production.

A careful analysis of the IPI fluctuations during moving states also revealed that fish drastically change the way they acquire sensory information when they perform certain behaviours. I find that IPIs become oscillatory when fish engage in rhythmic movement patterns around attended stimuli, and that two different types of frequency content exist. Although more work needs to be done to characterize it fully, this phenomenon could potentially serve as a window into the impact that rhythmic attentional processes can have on sensory sampling.

Finally, in Chapter 5, I model the electric fields generated by two interacting wave-type fish in order to uncover the sensory signature of endogenous vs. exogenous communication signals. How an animal resolves the problem of whether the environment has moved or it has moved, or whether a stimulus is generated by the environment or by itself, is a fundamental question in neuroscience. I find that the sensory signature of chirps is ambiguous, as multiple scenarios have the same impact on the AM pattern sensed by the fish. This novel finding, which breaks a traditionally held belief, raises the possibility that other cues are used by the animal to resolve the ambiguous states. Movement including interactions at different relative orientations likely are part of this

solution. It will take clever experiments to verify any such predictions.

To conclude, the contributions presented in this thesis add to a line of research [54] that serves as an alternative approach to the more common bottom-up, mechanistic and high-dimensional models of biological systems. I believe that the set of tools provided by this approach will catalyze the search for dynamical interpretations neural data, especially given the ongoing effort to reduce the dimensionality of experimental data sets [31]. Taken to its full potential, the search for simpler dynamical descriptions of neural systems, while still grounded in our understanding of microscopic dynamics, can equip us with the macroscopic language required to formulate general laws for the functioning of the brain.

References

- [1] P. A. Aguilera, A. C. Pereira, and Á. A. Caputi. Active electrolocation in pulse gymnotids: sensory consequences of objects' mutual polarization. *J Exp Biol*, 215(9):1533–1541, 2012. [28](#)
- [2] Y. Aït-Sahalia and L. Mancini. Out of sample forecasts of quadratic variation. *J Econom*, 147(1):17–33, 2008. [90](#), [91](#)
- [3] M. Anvari, M. R. R. Tabar, J. Peinke, and K. Lehnertz. Disentangling the stochastic behavior of complex time series. *Sci Rep*, 6:35435, 2016. [26](#), [75](#), [77](#), [105](#), [106](#)
- [4] G. Ashida and M. Kubo. Suprathreshold stochastic resonance induced by ion channel fluctuation. *Physica D*, 239(6):327–334, 2010. [108](#)
- [5] D. Babineau, A. Longtin, and J. E. Lewis. Modeling the electric field of weakly electric fish. *J Exp Biol*, 209(18):3636–3651, 2006. [131](#)
- [6] A. Bahraminasab, F. Ghasemi, A. Stefanovska, P. V. E. McClintock, and R. Friedrich. Physics of brain dynamics: Fokker–Planck analysis reveals changes in EEG δ – θ interactions in anaesthesia. *New J Phys*, 11(10):103051, 2009. [104](#)
- [7] R. S. Bains, S. Wells, R. R. Sillito, J. D. Armstrong, H. L. Cater, G. Banks, and P. M. Nolan. Assessing mouse behaviour throughout the light/dark cycle using automated in-cage analysis tools. *J Neurosci Meth*, 300:37–47, 2018. [138](#)
- [8] F. M. Bandi and T. H. Nguyen. On the functional estimation of jump–diffusion models. *J Econom*, 116(1):293–328, 2003. [75](#)
- [9] F. Bartolomei, P. Chauvel, and F. Wendling. Epileptogenicity of brain structures in human temporal lobe epilepsy: a quantified study from intracerebral EEG. *Brain*, 131(Pt 7):1818–1830, 2008. [25](#)

- [10] J. Bastian. Pyramidal-cell plasticity in weakly electric fish: a mechanism for attenuating responses to reafferent electrosensory inputs. *J. Comp. Physiol. A*, 176(1):63–73, 1995. [123](#)
- [11] S. Bazazi, F. Bartumeus, J. Hale, and I. Couzin. Intermittent motion in desert locusts: Behavioural complexity in simple environments. *PLoS Comput Biol*, 8(5):e1002498, 2012. [31](#), [71](#)
- [12] H. J. Bellen, C. Tong, and H. Tsuda. 100 years of *Drosophila* research and its impact on vertebrate neuroscience: a history lesson for the future. *Nat Rev Neurosci*, 11(7):514–522, 2010. [27](#)
- [13] J. Benda, A. Longtin, and L. Maler. Spike-frequency adaptation separates transient communication signals from background oscillations. *J. Neurosci.*, 25(9):2312–2321, 2005. [130](#), [131](#)
- [14] G. J. Berman, D. M. Choi, W. Bialek, and J. W. Shaevitz. Mapping the stereotyped behaviour of freely moving fruit flies. *J R Soc Interface*, 11(99):20140672, 2014. [25](#), [31](#), [138](#)
- [15] G. J. Berman, W. Bialek, and J. W. Shaevitz. Predictability and hierarchy in *Drosophila* behavior. *PNAS*, 113(42):11943–11948, 2016. [25](#)
- [16] N. J. Berman and L. Maler. Neural architecture of the electrosensory lateral line lobe: adaptations for coincidence detection, a sensory searchlight and frequency-dependent adaptive filtering. *J Exp Biol*, 202(10):1243–1253, 1999. [139](#)
- [17] D. Biswas, L. A. Arend, S. A. Stamper, B. P. Vágvolgyi, E. S. Fortune, and N. J. Cowan. Closed-Loop control of active sensing movements regulates sensory slip. *Current Biology*, 2018. [125](#)
- [18] K. Bol, G. Marsat, E. Harvey-Girard, A. Longtin, and L. Maler. Frequency-tuned cerebellar channels and burst-induced LTD lead to the cancellation of redundant sensory inputs. *J. Neurosci.*, 31(30):11028–11038, 2011. [108](#)
- [19] C. L. Buckley and T. Toyoizumi. A theory of how active behavior stabilises neural activity: Neural gain modulation by closed-loop environmental feedback. *PLoS Comput Biol*, 14(1):e1005926, 2018. [121](#)

- [20] T. H. Bullock, C. D. Hopkins, A. N. Popper, and R. R. Fay, editors. *Electroreception*, volume 21 of *Springer Handbook of Auditory Research*. Springer New York, 2005. [28](#), [29](#), [34](#)
- [21] A. N. Burkitt. A review of the integrate-and-fire neuron model: I. Homogeneous synaptic input. *Biol Cybern*, 95(1):1–19, 2006. [106](#)
- [22] D. J. Calderone, P. Lakatos, P. D. Butler, and F. X. Castellanos. Entrainment of neural oscillations as a modifiable substrate of attention. *Trends in Cognitive Sciences*, 18(6):300–309, 2014. [112](#), [124](#)
- [23] J. P. Capitanio and M. E. Emborg. Contributions of non-human primates to neuroscience research. *Lancet*, 371(9618):1126–1135, 2008. [27](#)
- [24] A. Capurro, A. Longtin, E. Bagarinao, S. Sato, O. Macadar, and K. Pakdaman. Variability of the electric organ discharge interval duration in resting *Gymnotus carapo*. *Biol Cybern*, 84(4):309–321, 2001. [27](#)
- [25] A. A. Caputi. The bioinspiring potential of weakly electric fish. *Bioinspir. Biomim.*, 12(2):025004, 2017. [27](#), [28](#)
- [26] A. A. Caputi, R. Budelli, K. Grant, and C. C. Bell. The electric image in weakly electric fish: physical images of resistive objects in *Gnathonemus petersii*. *J. Exp. Biol.*, 201(Pt 14):2115–2128, 1998. [28](#)
- [27] A. A. Caputi, P. A. Aguilera, and M. E. Castelló. Probability and amplitude of novelty responses as a function of the change in contrast of the reafferent image in *G. carapo*. *J Exp Biol*, 206(6):999–1010, 2003. [34](#)
- [28] C. Chow and J. White. Spontaneous action potentials due to channel fluctuations. *Biophys J*, 71:3013–21, 1997. [106](#)
- [29] P. T. Clemson and A. Stefanovska. Discerning non-autonomous dynamics. *Phys Rep*, 542(4):297–368, 2014. [74](#)
- [30] J. J. Collins, C. C. Chow, A. C. Capela, and T. T. Imhoff. Aperiodic stochastic resonance. *Phys Rev E*, 54(5):5575–5584, 1996. [66](#)
- [31] J. P. Cunningham and B. M. Yu. Dimensionality reduction for large-scale neural recordings. *Nat Neurosci*, 17(11):1500–1509, 2014. [141](#)

- [32] J. C. Curtis and D. Kleinfeld. Phase-to-rate transformations encode touch in cortical neurons of a scanning sensorimotor system. *Nat Neurosci*, 12(4):492–501, 2009. [123](#)
- [33] C. Curto, S. Sakata, S. Marguet, V. Itskov, and K. D. Harris. A simple model of cortical dynamics explains variability and state dependence of sensory responses in urethane-anesthetized auditory cortex. *J Neurosci*, 29(34):10600–10612, 2009. [33](#)
- [34] T. Czernik, J. Kula, J. Łuczka, and P. Hanggi. Thermal ratchets driven by Poissonian white shot noise. *Phys. Rev. E*, 55, 1997. [75](#)
- [35] E. Daly and A. Porporato. Probabilistic dynamics of some jump-diffusion systems. *Phys Rev E*, 73(2):026108, 2006. [75](#)
- [36] P. Dayan and L. F. Abbott. *Theoretical Neuroscience: Computational and Mathematical Modeling of Neural Systems*. The MIT Press, Cambridge, Mass., 1 edition edition, 2005. [8](#)
- [37] G. Deco, D. Martí, A. Ledberg, R. Reig, and M. V. Sanchez Vives. Effective reduced diffusion-models: a data driven approach to the analysis of neuronal dynamics. *PLoS Comput Biol*, 5(12):e1000587, 2009. [33](#), [74](#)
- [38] M. Deschênes, J. Moore, and D. Kleinfeld. Sniffing and whisking in rodents. *Curr. Opin. Neurobiol.*, 22(2):243–250, 2012. [123](#), [124](#)
- [39] K. Diba, H. A. Lester, and C. Koch. Intrinsic noise in cultured hippocampal neurons: experiment and modeling. *J. Neurosci.*, 24(43):9723–9733, 2004. [104](#)
- [40] A. D. Dorval and J. A. White. Channel noise is essential for perithreshold oscillations in entorhinal stellate neurons. *J. Neurosci.*, 25(43):10025–10028, 2005. [103](#), [106](#)
- [41] T. W. Dunn, Y. Mu, S. Narayan, O. Randlett, E. A. Naumann, C.-T. Yang, A. F. Schier, J. Freeman, F. Engert, and M. B. Ahrens. Brain-wide mapping of neural activity controlling zebrafish exploratory locomotion. *eLife*, 5:e12741, 2016. [138](#)
- [42] B. Ellenbroek and J. Youn. Rodent models in neuroscience research: is it a rat race? *Dis Model Mech*, 9(10):1079–1087, 2016. [27](#)

- [43] S. B. Elliott and L. Maler. Stimulus induced Up states in the dorsal pallium of a weakly electric fish. *J Neurophysiol*, 114(3):2071–6, 2015. [68](#), [72](#)
- [44] G. v. d. Emde and T. Ringer. Electrolocation of capacitive objects in four species of pulse-type weakly electric fish I. discrimination performance. *Ethology*, 91(4):326–338, 1992. [112](#)
- [45] R. F. Pawula. Approximation of the linear boltzmann equation by the fokker-planck equation. *Phys Rev*, 162:186–188, 1967. [26](#), [77](#)
- [46] A. A. Faisal and S. B. Laughlin. Stochastic simulations on the reliability of action potential propagation in thin axons. *PLoS Comput Biol*, 3(5):e79, 2007. [106](#)
- [47] A. A. Faisal, L. P. J. Selen, and D. M. Wolpert. Noise in the nervous system. *Nat Rev Neurosci*, 9(4):292–303, 2008. [106](#)
- [48] S. Fauve and F. Heslot. Stochastic resonance in a bistable system. *Phys Lett A*, 97(1–2):5–7, 1983. [43](#)
- [49] A. Feldmann and W. Whitt. Fitting mixtures of exponentials to long-tail distributions to analyze network performance models. *Performance Evaluation*, 31(3–4):245–279, 1998. [71](#)
- [50] I. C. Fiebelkorn, M. A. Pinsk, and S. Kastner. A dynamic interplay within the frontoparietal network underlies rhythmic spatial attention. *Neuron*, 99(4):842–853.e8, 2018. [124](#)
- [51] J. E. Figueroa-López and J. Nisen. Optimally thresholded realized power variations for Lévy jump diffusion models. *Stoch Process Their Appl*, 123(7):2648–2677, 2013. [75](#)
- [52] R. Fox and Y. Lu. Emergent collective behavior in large numbers of globally coupled independently stochastic ion channels. *Phys Rev E*, 49(4):3421–3431, 1994. [107](#)
- [53] R. F. Fox. Stochastic versions of the Hodgkin-Huxley equations. *Biophys J*, 72(5):2068–2074, 1997. [107](#)
- [54] R. Friedrich, J. Peinke, M. Sahimi, and M. Reza Rahimi Tabar. Approaching complexity by stochastic methods: From biological systems to turbulence. *Phys Rep*, 506(5):87–162, 2011. [26](#), [33](#), [74](#), [77](#), [104](#), [141](#)

- [55] L. Gammaitoni, P. Hänggi, P. Jung, and F. Marchesoni. Stochastic resonance. *Rev Mod Phys*, 70(1):223–287, 1998. [71](#)
- [56] L. Gao, S. Balakrishnan, W. He, Z. Yan, and R. Müller. Ear Deformations Give Bats a Physical Mechanism for Fast Adaptation of Ultrasonic Beam Patterns. *Phys. Rev. Lett.*, 107(21):214301, 2011. [113](#)
- [57] C. W. Gardiner. *Stochastic methods: a handbook for the natural and social sciences*. Springer series in synergetics. Springer, Berlin, 4th ed edition, 2009. [20](#), [93](#), [109](#), [126](#)
- [58] W. Geng, P. Cosman, J.-H. Baek, C. C. Berry, and W. R. Schafer. Quantitative classification and natural clustering of *Caenorhabditis elegans* behavioral phenotypes. *Genetics*, 165(3):1117–1126, 2003. [23](#)
- [59] M. Ghorbani, M. Mehta, R. Bruinsma, and A. J. Levine. Nonlinear-dynamics theory of up-down transitions in neocortical neural networks. *Phys Rev E*, 85(2):021908, 2012. [69](#)
- [60] A. C. Giassi, E. Harvey-Girard, B. Valsamis, and L. Maler. Organization of the gymnotiform fish pallium in relation to learning and memory: I. Cytoarchitectonics and cellular morphology. *J Comp Neurol*, 520(15):3314–3337, 2012. [34](#), [49](#), [68](#)
- [61] D. T. Gillespie. *Markov Processes: An Introduction for Physical Scientists*. Academic Press, Boston, 1 edition edition, 1991. [17](#), [18](#), [82](#), [109](#)
- [62] D. T. Gillespie and E. Seitaridou. *Simple Brownian diffusion: an introduction to the standard theoretical models*. Oxford University Press, Oxford, United Kingdom, first edition edition, 2013. [10](#), [13](#), [14](#), [94](#), [105](#), [109](#)
- [63] M. T. Giraudo and L. Sacerdote. Jump-diffusion processes as models for neuronal activity. *Biosystems*, 40(1):75–82, 1997. [75](#)
- [64] M. T. Giraudo, L. Sacerdote, and R. Sirovich. Effects of random jumps on a very simple neuronal diffusion model. *Biosystems*, 67(1):75–83, 2002. [75](#)
- [65] J. H. Goldwyn and E. Shea-Brown. The what and where of adding channel noise to the Hodgkin-Huxley equations. *PLoS Comput Biol*, 7(11):e1002247, 2011. [106](#), [107](#)

- [66] J. H. Goldwyn, N. S. Imennov, M. Famulare, and E. Shea-Brown. Stochastic differential equation models for ion channel noise in Hodgkin-Huxley neurons. *Phys Rev E*, 83(4 Pt 1):041908, 2011. [107](#)
- [67] S. Grillner, J. Hellgren, A. Ménard, K. Saitoh, and M. A. Wikström. Mechanisms for selection of basic motor programs – roles for the striatum and pallidum. *Trends Neurosci*, 28(7):364–370, 2005. [68](#)
- [68] S. Haegens and E. Zion Golumbic. Rhythmic facilitation of sensory processing: A critical review. *Neuroscience & Biobehavioral Reviews*, 86:150–165, 2018. [111](#)
- [69] H. Haken. *Information and self-organization: a macroscopic approach to complex systems*. Springer-Verlag, Berlin; New York, 1988. OCLC: 649906343. [74](#)
- [70] P. Hansen and A. Lunde. A Realized Variance for the Whole Day Based on Intermittent High-Frequency Data. *J financ Econom*, 3:525–554, 2005. [90](#), [91](#)
- [71] A. Harnos, G. Horváth, A. B. Lawrence, and G. Vattay. Scaling and intermittency in animal behaviour. *Physica A*, 286(1–2):312–320, 2000. [27](#), [71](#)
- [72] E. Harvey-Girard, A. C. C. Giassi, W. Ellis, and L. Maler. Expression of the cannabinoid CB1 receptor in the gymnotiform fish brain and its implications for the organization of the teleost pallium. *J Comp Neurol*, 521(4):949–975, 2013. [68](#)
- [73] W. Heiligenberg. The jamming avoidance response of the electric fish, Eigenmannia: computational rules and their neuronal implementation. *Seminars in Neuroscience*, 3(1):3–18, 1991. [129](#)
- [74] R. F. Helfrich, I. C. Fiebelkorn, S. M. Szczepanski, J. J. Lin, J. Parvizi, R. T. Knight, and S. Kastner. Neural mechanisms of sustained attention are rhythmic. *Neuron*, 99(4):854–865.e5, 2018. [124](#)
- [75] J. Henninger, R. Krahe, F. Kirschbaum, J. Grewe, and J. Benda. Statistics of natural communication signals observed in the wild identify important yet neglected stimulus regimes in weakly electric fish. *J. Neurosci.*, 38(24):5456–5465, 2018. [130](#)
- [76] J. Hidalgo, L. F. Seoane, J. M. Cortés, and M. A. Muñoz. Stochastic amplification of fluctuations in cortical up-states. *PLoS ONE*, 7(8):e40710, 2012. [69](#)

- [77] C. Hillar, G. Omnis, D. Rhea, and L. Tecott. Active state organization of spontaneous behavioral patterns. *Sci Rep*, 8(1):1064, 2018. [138](#)
- [78] R. Hindriks, F. Bijma, B. W. van Dijk, Y. D. van der Werf, E. J. W. van Someren, and A. W. van der Vaart. Dynamics underlying spontaneous human alpha oscillations: A data-driven approach. *NeuroImage*, 57(2):440–451, 2011. [26](#), [33](#), [74](#)
- [79] R. Hindriks, R. Jansen, F. Bijma, H. D. Mansvelder, M. C. M. de Gunst, and A. W. van der Vaart. Unbiased estimation of Langevin dynamics from time series with application to hippocampal field potentials in vitro. *Phys Rev E*, 84(2):021133, 2011. [26](#), [33](#), [74](#)
- [80] V. Hofmann, J. I. Sanguinetti-Scheck, S. Künzel, B. Geurten, L. Gómez-Sena, and J. Engelmann. Sensory flow shaped by active sensing: sensorimotor strategies in electric fish. *J Exp Biol*, 216(13):2487–2500, 2013. [111](#), [112](#)
- [81] V. Hofmann, B. R. H. Geurten, J. I. Sanguinetti-Scheck, L. Gómez-Sena, and J. Engelmann. Motor patterns during active electrosensory acquisition. *Front. Behav. Neurosci.*, 8:186, 2014. [27](#), [112](#)
- [82] V. Hofmann, J. I. Sanguinetti-Scheck, L. Gómez-Sena, and J. Engelmann. Sensory flow as a basis for a novel distance cue in freely behaving electric fish. *J. Neurosci.*, pages 1361–16, 2016. [111](#)
- [83] D. Holcman and M. Tsodyks. The emergence of up and down states in cortical networks. *PLoS Computat Biol*, 2(3):e23, 2006. [69](#)
- [84] Y. Holovatch, R. Kenna, and S. Thurner. Complex systems: physics beyond physics. *Eur. J. Phys.*, 38(2):023002, 2017. [74](#)
- [85] H. P. Hsu. *Schaum’s Outline of Probability, Random Variables, and Random Processes, 3rd Edition*. McGraw-Hill Education, New York, 3 edition edition, 2014. [10](#), [110](#)
- [86] G. Huesa, A. N. van den Pol, and T. E. Finger. Differential distribution of hypocretin (orexin) and melanin-concentrating hormone in the goldfish brain. *J Comp Neurol*, 488(4):476–491, 2005. [68](#)

- [87] G. J. Hupé and J. E. Lewis. Electrocommunication signals in free swimming brown ghost knifefish, *Apteronotus leptorhynchus*. *Journal of Experimental Biology*, 211(10):1657–1667, 2008. [28](#), [130](#)
- [88] G. J. Hupé, J. E. Lewis, and J. Benda. The effect of difference frequency on electrocommunication: chirp production and encoding in a species of weakly electric fish, *Apteronotus leptorhynchus*. *J. Physiol. Paris*, 102(4-6):164–172, 2008. [28](#), [130](#)
- [89] G. A. Jacobson, K. Diba, A. Yaron-Jakoubovitch, Y. Oz, C. Koch, I. Segev, and Y. Yarom. Subthreshold voltage noise of rat neocortical pyramidal neurones. *J Physiol*, 564(Pt 1):145–160, 2005. [107](#)
- [90] K. L. Jens Prusseit. Stochastic qualifiers of epileptic brain dynamics. *Phys Rev Lett*, 98(13):138103, 2007. [33](#)
- [91] M. Johannes. The Statistical and Economic Role of Jumps in Continuous-Time Interest Rate Models. *J Finance*, 59(1):227–260, 2004. [75](#)
- [92] D. C. Johnston. Stretched exponential relaxation arising from a continuous sum of exponential decays. *Phys Rev B*, 74(18):184430, 2006. [71](#)
- [93] G. Jones and M. W. Holderied. Bat echolocation calls: adaptation and convergent evolution. *Proc Biol Sci*, 274(1612):905–912, 2007. [113](#)
- [94] J. J. Jun, A. Longtin, and L. Maler. Precision measurement of electric organ discharge timing from freely moving weakly electric fish. *J Neurophysiol*, 107(7):1996–2007, 2012. [36](#), [113](#)
- [95] J. J. Jun, A. Longtin, and L. Maler. Enhanced sensory sampling precedes self-initiated locomotion in an electric fish. *J Exp Biol*, 217(20):3615–3628, 2014. [32](#), [34](#), [35](#), [36](#), [37](#), [49](#), [67](#), [69](#), [113](#)
- [96] J. J. Jun, A. Longtin, and L. Maler. Long-term behavioral tracking of freely swimming weakly electric fish. *J Vis Exp*, (85), 2014. [36](#), [113](#), [126](#)
- [97] J. J. Jun, A. Longtin, and L. Maler. Active sensing associated with spatial learning reveals memory-based attention in an electric fish. *J Neurophysiol*, 115(5):2577–2592, 2016. [27](#), [80](#), [108](#), [112](#), [115](#), [124](#)

- [98] K. Kagaya and M. Takahata. Readiness Discharge for Spontaneous Initiation of Walking in Crayfish. *J Neurosci*, 30(4):1348–1362, 2010. [31](#)
- [99] A. V. Kalueff, A. M. Stewart, and R. Gerlai. Zebrafish as an emerging model for studying complex brain disorders. *Trends Pharmacol Sci*, 35(2):63–75, 2014. [27](#)
- [100] J. Kaslin, J. M. Nystedt, M. Ostergård, N. Peitsaro, and P. Panula. The orexin/hypocretin system in zebrafish is connected to the aminergic and cholinergic systems. *J Neurosci*, 24(11):2678–2689, 2004. [68](#)
- [101] A. K. Katsaggelos. Iterative Image Restoration Algorithms. *Opt Eng*, 28(7):287735, 1989. [89](#)
- [102] J. Keifer and C. H. Summers. Putting the “Biology” Back into “Neurobiology”: The Strength of Diversity in Animal Model Systems for Neuroscience Research. *Front Syst Neurosci*, 10, 2016. [26](#)
- [103] M. Kelly, D. Babineau, A. Longtin, and J. E. Lewis. Electric field interactions in pairs of electric fish: modeling and mimicking naturalistic inputs. *Biol Cybern*, 98(6):479–490, 2008. [131](#), [132](#)
- [104] C. Koch. *Biophysics of Computation: Information Processing in Single Neurons*. Oxford University Press, New York, 2005. [8](#)
- [105] H. H. Kornhuber and L. Deecke. Hirnpotentialänderungen bei Willkürbewegungen und passiven Bewegungen des Menschen: Bereitschaftspotential und reafferente Potentiale. *Pflüger’s Archiv für die gesamte Physiologie des Menschen und der Tiere*, 284(1):1–17, 1965. [70](#)
- [106] S. G. Kou. Jump-Diffusion Models for Asset Pricing in Financial Engineering. In J. R. Birge and V. Linetsky, editors, *Handbooks in Operations Research and Management Science*, volume 15 of *Financial Engineering*, pages 73–116. Elsevier, 2007. doi: 10.1016/S0927-0507(07)15002-7. [75](#)
- [107] P. Kounitsky, J. Rydell, E. Amichai, A. Boonman, O. Eitan, A. J. Weiss, and Y. Yovel. Bats adjust their mouth gape to zoom their biosonar field of view. *PNAS*, page 201422843, 2015. [113](#)
- [108] D. L. Kramer and R. L. McLaughlin. The behavioral ecology of intermittent locomotion. *Am Zool*, 41(2):137–153, 2001. [27](#), [31](#), [33](#)

- [109] P. Lakatos, G. Karmos, A. D. Mehta, I. Ulbert, and C. E. Schroeder. Entrainment of neuronal oscillations as a mechanism of attentional selection. *Science*, 320(5872):110–113, 2008. [112](#), [124](#)
- [110] D. Lamouroux and K. Lehnertz. Kernel-based regression of drift and diffusion coefficients of stochastic processes. *Phys Lett A*, 373(39):3507–3512, 2009. [33](#), [74](#)
- [111] E. W. Large and M. R. Jones. The dynamics of attending: How people track time-varying events. *Psych Rev*, 106(1):119–159, 1999. [112](#)
- [112] S. D. Larson, P. Gleeson, and A. E. X. Brown. Connectome to behaviour: modelling *Caenorhabditis elegans* at cellular resolution. *Philos Trans R Soc Lond B*, 373(1758), 2018. [23](#)
- [113] S. S. Lee and J. Hannig. Detecting jumps from Lévy jump diffusion processes. *J Financ Econ*, 96(2):271–290, 2010. [75](#)
- [114] P. N. Lehner. *Handbook of Ethological Methods*. Cambridge University Press, Cambridge, 2 edition edition, 1998. [23](#)
- [115] M. S. Lewicki. A review of methods for spike sorting: the detection and classification of neural action potentials. *Network*, 9(4):R53–R78, 1998. [38](#)
- [116] B. Libet, C. A. Gleason, E. W. Wright, and D. K. Pearl. Time of conscious intention to act in relation to onset of cerebral activity (readiness-potential). The unconscious initiation of a freely voluntary act. *Brain*, 106 (Pt 3):623–642, 1983. [33](#), [70](#)
- [117] L. S. Liebovitch and T. I. Tóth. Distributions of activation energy barriers that produce stretched exponential probability distributions for the time spent in each state of the two state reaction $A \rightleftharpoons B$. *Bull Math Biol*, 53(3):443–455, 1991. [71](#)
- [118] D. Linaro, M. Storace, and M. Giugliano. Accurate and Fast Simulation of Channel Noise in Conductance-Based Model Neurons by Diffusion Approximation. *PLoS Comput Biol*, 7(3):e1001102, 2011. [107](#)
- [119] F. Lopes da Silva. EEG and MEG: Relevance to Neuroscience. *Neuron*, 80(5):1112–1128, 2013. [25](#)

- [120] J. Luczka, R. Bartussek, and P. Hänggi. White-noise-induced transport in periodic structures. *Europhys Lett*, 31(8):431, 1995. [75](#)
- [121] J. Łuczka, T. Czernik, and P. Hanggi. Symmetric white noise can induce directed current in ratchets. *Phys. Rev. E*, 56, 1997. [75](#)
- [122] J.-R. Luevano. Statistical features of the Stretched Exponentials Densities. *J Phys Conf Ser*, 475:012008, 2013. [62](#), [66](#)
- [123] A. Maesani, P. Ramdya, S. Cruchet, K. Gustafson, R. Benton, and D. Floreano. Fluctuation-Driven Neural Dynamics Reproduce Drosophila Locomotor Patterns. *PLOS Comput Biol*, 11(11):e1004577, 2015. [25](#), [68](#)
- [124] C. M. Marcoux, S. E. Clarke, W. H. Nesse, A. Longtin, and L. Maler. Balanced ionotropic receptor dynamics support signal estimation via voltage-dependent membrane noise. *J Neurophysiol*, 115(1):530–545, 2016. [99](#), [101](#), [103](#), [106](#), [107](#), [108](#)
- [125] M. R. Markham. Electrocyte physiology: 50 years later. *Journal of Experimental Biology*, 216(13):2451–2458, 2013. [27](#)
- [126] D. Martí, G. Deco, M. Mattia, G. Gigante, and P. Del Giudice. A fluctuation-driven mechanism for slow decision processes in reverberant networks. *PLoS ONE*, 3(7):e2534, 2008. [48](#)
- [127] A. Mashanova, T. H. Oliver, and V. A. A. Jansen. Evidence for intermittency and a truncated power law from highly resolved aphid movement data. *J R Soc Interface*, 7(42):199–208, 2010. [71](#)
- [128] A. Maye, C.-h. Hsieh, G. Sugihara, and B. Brembs. Order in spontaneous behavior. *PLoS ONE*, 2(5):e443, 2007. [25](#), [31](#)
- [129] M. D. McDonnell and L. M. Ward. The benefits of noise in neural systems: bridging theory and experiment. *Nat. Rev. Neurosci.*, 12(7):415–426, 2011. [106](#)
- [130] J. F. Mejias, H. J. Kappen, and J. J. Torres. Irregular dynamics in up and down cortical states. *PLoS ONE*, 5(11):e13651, 2010. [33](#), [69](#), [71](#)
- [131] A. Melanson and A. Longtin. Data-driven jump-diffusion modelling with application to membrane voltage fluctuations in pyramidal neurons. (*submitted*). [126](#)

- [132] A. Melanson, J. F. Mejias, J. J. Jun, L. Maler, and A. Longtin. Nonstationary stochastic dynamics underlie spontaneous transitions between active and inactive behavioral states. *eNeuro*, 4(2), 2017. [114](#), [126](#)
- [133] M. G. Metzen and M. J. Chacron. Weakly electric fish display behavioral responses to envelopes naturally occurring during movement: implications for neural processing. *J Exp Biol*, 217(8):1381–1391, 2014. [130](#)
- [134] M. G. Metzen and M. J. Chacron. Stimulus background influences phase invariant coding by correlated neural activity. *eLife*, 6:e24482, 2017. [130](#)
- [135] B. Mitchinson and T. J. Prescott. Whisker movements reveal spatial attention: a unified computational model of active sensing control in the rat. *PLoS Comput Biol*, 9(9):e1003236, 2013. [124](#)
- [136] K. T. Moortgat, T. H. Bullock, and T. J. Sejnowski. Precision of the pacemaker nucleus in a weakly electric fish: network versus cellular influences. *J. Neurophysiol.*, 83(2):971–983, 2000. [27](#)
- [137] B. Morillon and S. Baillet. Motor origin of temporal predictions in auditory attention. *PNAS*, 114(42):E8913–E8921, 2017. [124](#)
- [138] B. Morillon, C. E. Schroeder, and V. Wyart. Motor contributions to the temporal precision of auditory attention. *Nat Comm*, 5:5255, 2014. [124](#)
- [139] B. Morillon, T. A. Hackett, Y. Kajikawa, and C. E. Schroeder. Predictive motor control of sensory dynamics in auditory active sensing. *Curr Opin Neurobiol*, 31:230–238, 2015. [124](#)
- [140] S. M. Mousavi, S. N. S. Reihani, G. Anvari, M. Anvari, H. G. Alinezhad, and M. R. R. Tabar. Stochastic analysis of time series for the spatial positions of particles trapped in optical tweezers. *Sci Rep*, 7(1):4832, 2017. [74](#)
- [141] M. Murakami, M. I. Vicente, G. M. Costa, and Z. F. Mainen. Neural antecedents of self-initiated actions in secondary motor cortex. *Nat Neurosci*, 17(11):1574–1582, 2014. [33](#)
- [142] M. Musila and P. Lánský. Generalized Stein’s model for anatomically complex neurons. *Biosystems*, 25(3):179–191, 1991. [75](#)

- [143] M. E. Nelson. Electric fish. *Current Biology*, 21(14):R528–R529, 2011. [27](#)
- [144] M. E. Nelson, Z. Xu, and J. R. Payne. Characterization and modeling of P-type electrosensory afferent responses to amplitude modulations in a wave-type electric fish. *J Comp Physiol A*, 181(5):532–544, 1997. [139](#)
- [145] G. Nicolis and C. Rouvas-Nicolis. Complex systems. *Scholarpedia*, 2(11):1473, 2007. [74](#)
- [146] L. Noonan, B. Doiron, C. Laing, A. Longtin, and R. W. Turner. A dynamic dendritic refractory period regulates burst discharge in the electrosensory lobe of weakly electric fish. *J. Neurosci.*, 23(4):1524–1534, 2003. [99](#)
- [147] P. Orio and D. Soudry. Simple, fast and accurate implementation of the diffusion approximation algorithm for stochastic ion channels with multiple states. *PLOS ONE*, 7(5):e36670, 2012. [107](#)
- [148] T. Palmér, F. Ek, O. Enqvist, R. Olsson, K. Åström, and P. Petersson. Action sequencing in the spontaneous swimming behavior of zebrafish larvae - implications for drug development. *Sci Rep*, 7, 2017. [138](#)
- [149] A. Papoulis and S. U. Pillai. *Probability, Random Variables and Stochastic Processes with Errata Sheet*. McGraw-Hill Higher Education, Boston, Mass., 4 edition, 2002. [84](#), [110](#)
- [150] S. V. Patankar. *Numerical heat transfer and fluid flow*. Series in computational methods in mechanics and thermal sciences. Hemisphere Publ. Co, New York, 1980. OCLC: 31743097. [21](#), [22](#), [61](#), [82](#), [109](#)
- [151] R. K. Pathria and P. D. Beale. *Statistical mechanics*. Elsevier/Academic Press, Amsterdam ; Boston, 3rd ed edition, 2011. [16](#), [109](#)
- [152] A. C. Pereira, A. Rodríguez-Cattáneo, and A. A. Caputi. The slow pathway in the electrosensory lobe of *Gymnotus omarorum*: Field potentials and unitary activity. *J Physiol Paris*, 108(2–3):71–83, 2014. [68](#)
- [153] M. Petelczyc, J. J. Zebrowski, and R. Baranowski. Kramers-Moyal coefficients in the analysis and modeling of heart rate variability. *Phys Rev E*, 80:031127, 2009. [26](#), [74](#)

- [154] M. Petelczyc, J. J. Żebrowski, and E. Orłowska-Baranowska. A fixed mass method for the Kramers-Moyal expansion—Application to time series with outliers. *Chaos*, 25(3):033115, 2015. [26](#), [74](#)
- [155] A. S. Pikovsky and J. Kurths. Coherence resonance in a noise-driven excitable system. *Phys Rev Lett*, 78(5):775–778, 1997. [70](#), [120](#)
- [156] T. J. Prescott, M. E. Diamond, and A. M. Wing. Active touch sensing. *Philos Trans R Soc Lond B Biol Sci*, 366(1581):2989–2995, 2011. [111](#)
- [157] D. A. Prober, J. Rihel, A. A. Onah, R.-J. Sung, and A. F. Schier. Hypocretin/orexin overexpression induces an insomnia-like phenotype in zebrafish. *J Neurosci*, 26(51):13400–13410, 2006. [49](#), [68](#)
- [158] A. Proekt, J. R. Banavar, A. Maritan, and D. W. Pfaff. Scale invariance in the dynamics of spontaneous behavior. *Proc Natl Acad Sci U S A*, 109(26):10564–10569, 2012. [27](#), [31](#), [71](#)
- [159] J. Prusseit and K. Lehnertz. Stochastic qualifiers of epileptic brain dynamics. *Phys. Rev. Lett.*, 98(13):138103, 2007. [25](#), [74](#)
- [160] M. Ragwitz and H. Kantz. Indispensable finite time corrections for Fokker-Planck equations from time series data. *Phys Rev Lett*, 87(25):254501, 2001. [74](#)
- [161] C. Rajkai, P. Lakatos, C.-M. Chen, Z. Pincze, G. Karmos, and C. E. Schroeder. Transient cortical excitation at the onset of visual fixation. *Cereb. Cortex*, 18(1):200–209, 2008. [123](#)
- [162] B. Rasnow. The effects of simple objects on the electric field of Apterionotus. *J Comp Physiol A*, 178(3):397–411, 1996. [28](#)
- [163] M. J. E. Richardson and W. Gerstner. Statistics of subthreshold neuronal voltage fluctuations due to conductance-based synaptic shot noise. *Chaos*, 16(2):026106, 2006. [107](#)
- [164] D. L. Ringach. Spontaneous and driven cortical activity: implications for computation. *Curr. Opin. Neurobiol.*, 19(4):439–444, 2009. [139](#)
- [165] H. Risken and T. Frank. *The Fokker-Planck Equation: Methods of Solution and Applications*. Springer Series in Synergetics. Springer-Verlag, Berlin Heidelberg, 2 edition, 1996. [82](#), [94](#)

- [166] P. Rowat. Interspike interval statistics in the stochastic Hodgkin-Huxley model: coexistence of gamma frequency bursts and highly irregular firing. *Neural Comput*, 19(5):1215–1250, 2007. [107](#)
- [167] A. Roxin and A. Ledberg. Neurobiological models of two-choice decision making can be reduced to a one-dimensional nonlinear diffusion equation. *PLoS Comput Biol*, 4(3):e1000046, 2008. [49](#)
- [168] M. Rudolph and A. Destexhe. Correlation detection and resonance in neural systems with distributed noise sources. *Phys. Rev. Lett.*, 86(16):3662–3665, 2001. [108](#)
- [169] B. A. Schmerl and M. D. McDonnell. Channel-noise-induced stochastic facilitation in an auditory brainstem neuron model. *Phys. Rev. E*, 88(5):052722, 2013. [106](#), [108](#)
- [170] E. Schneidman, B. Freedman, and I. Segev. Ion channel stochasticity may be critical in determining the reliability and precision of spike timing. *Neural Comput*, 10(7):1679–1703, 1998. [106](#)
- [171] C. E. Schroeder and P. Lakatos. Low-frequency neuronal oscillations as instruments of sensory selection. *Trends Neurosci.*, 32(1):9–18, 2009. [112](#), [124](#)
- [172] C. E. Schroeder, D. A. Wilson, T. Radman, H. Scharfman, and P. Lakatos. Dynamics of active sensing and perceptual selection. *Curr Opin Neurobiol*, 20(2):172–176, 2010. [111](#), [112](#), [123](#), [124](#)
- [173] A. Schurger, J. D. Sitt, and S. Dehaene. An accumulator model for spontaneous neural activity prior to self-initiated movement. *Proc Natl Acad Sci U S A*, 109(42):E2904–E2913, 2012. [33](#), [70](#)
- [174] M. N. Shadlen and R. Kiani. Decision making as a window on cognition. *Neuron*, 80(3):791–806, 2013. [70](#)
- [175] T. O. Sharpee. Taking a close look at electrosensing. *eLife*, 5:e16209, 2016. [130](#)
- [176] R. Shusterman, M. C. Smear, A. A. Koulakov, and D. Rinberg. Precise olfactory responses tile the sniff cycle. *Nat Neurosci*, 14(8):1039–1044, 2011. [123](#)

- [177] S. Siegert, R. Friedrich, and J. Peinke. Analysis of data sets of stochastic systems. *Phys Lett A*, 243(5):275–280, 1998. [74](#)
- [178] R. Sirovich and L. Sacerdote. Noise induced phenomena in jump diffusion models for single neuron spike activity. In *2004 IEEE International Joint Conference on Neural Networks (IEEE Cat. No.04CH37541)*, volume 4, pages 3025–3028, 2004. doi: 10.1109/IJCNN.2004.1381149. [75](#)
- [179] R. Sirovich, L. Sacerdote, and A. E. P. Villa. Cooperative behavior in a jump diffusion model for a simple network of spiking neurons. *Math Biosci Eng*, 11(2): 385–401, 2014. [75](#)
- [180] S. J. M. Smith. EEG in the diagnosis, classification, and management of patients with epilepsy. *Journal of Neurology, Neurosurgery & Psychiatry*, 76(suppl 2):ii2–ii7, 2005. [25](#)
- [181] N. Spruston. Pyramidal neuron. *Scholarpedia*, 4(5):6130, 2009. [28](#)
- [182] S. A. Stamper, M. S. Madhav, N. J. Cowan, and E. S. Fortune. Beyond the Jamming Avoidance Response: weakly electric fish respond to the envelope of social electrosensory signals. *J Exp Biol*, 215(23):4196–4207, 2012. [28](#)
- [183] S. A. Stamper, E. Roth, N. J. Cowan, and E. S. Fortune. Active sensing via movement shapes spatiotemporal patterns of sensory feedback. *J Exp Biol*, 215(9): 1567–1574, 2012. [125](#)
- [184] P. N. Steinmetz, A. Manwani, C. Koch, M. London, and I. Segev. Subthreshold voltage noise of rat neocortical pyramidal neurones. *J Comput Neurosci*, 9(2): 133–148, 2000. [107](#)
- [185] G. J. Stephens, B. Johnson-Kerner, W. Bialek, and W. S. Ryu. Dimensionality and dynamics in the behavior of *C. elegans*. *PLoS Comput Biol*, 4(4):e1000028, 2008. [23](#)
- [186] G. J. Stephens, M. B. d. Mesquita, W. S. Ryu, and W. Bialek. Emergence of long timescales and stereotyped behaviors in *Caenorhabditis elegans*. *PNAS*, 108(18): 7286–7289, 2011. [23](#)

- [187] G. J. Stephens, L. C. Osborne, and W. Bialek. Searching for simplicity in the analysis of neurons and behavior. *Proc Natl Acad Sci U S A*, 108:15565–15571, 2011. [32](#)
- [188] D. Sterratt, B. Graham, D. A. Gillies, and D. Willshaw. *Principles of Computational Modelling in Neuroscience*. Cambridge University Press, Cambridge ; New York, 1 edition edition, 2011. [7](#)
- [189] E. W. Tan, J. M. Nizar, E. Carrera-G, and E. S. Fortune. Electrosensory interference in naturally occurring aggregates of a species of weakly electric fish, *Eigenmannia virescens*. *Behav. Brain Res.*, 164(1):83–92, 2005. [129](#)
- [190] P. Tankov and E. Voltchkova. Jump-diffusion models: a practitioner’s guide. *Banque et Marchés*, 2009. [75](#)
- [191] M. J. Toerring and P. Belbenoit. Motor Programmes and Electroreception in Mormyrid Fish. *Behavioral Ecology and Sociobiology*, 4(4):369–379, 1979. [112](#)
- [192] A.-T. Trinh, E. Harvey-Girard, F. Teixeira, and L. Maler. Cryptic laminar and columnar organization in the dorsolateral pallium of a weakly electric fish. *The Journal of Comparative Neurology*, 524(2):408–428, 2016. [68](#)
- [193] Y. Tu and G. Grinstein. How white noise generates power-law switching in bacterial flagellar motors. *Phys Rev Lett*, 94(20):208101, 2005. [71](#)
- [194] R. W. Turner, L. Maler, T. Deerinck, S. R. Levinson, and M. H. Ellisman. TTX-sensitive dendritic sodium channels underlie oscillatory discharge in a vertebrate sensory neuron. *J. Neurosci.*, 14(11):6453–6471, 1994. [107](#)
- [195] M. Tutkun and L. Mydlarski. Markovian properties of passive scalar increments in grid-generated turbulence. *New J Phys*, 6(1):49, 2004. [74](#)
- [196] A. M. van Mourik, A. Daffertshofer, and P. J. Beek. Estimating Kramers–Moyal coefficients in short and non-stationary data sets. *Phys Lett A*, 351(1):13–17, 2006. [74](#)
- [197] M. Wachowiak. All in a sniff: olfaction as a model for active sensing. *Neuron*, 71(6):962–973, 2011. [111](#), [124](#)

- [198] X.-J. Wang. Probabilistic decision making by slow reverberation in cortical circuits. *Neuron*, 36(5):955–968, 2002. [48](#)
- [199] X.-J. Wang. Decision making in recurrent neuronal circuits. *Neuron*, 60(2):215–234, 2008. [70](#)
- [200] J. White, E. Southgate, J. Thomson, and S. Brenner. The Structure of the Nervous System of the Nematode *Caenorhabditis elegans*. *Philos Trans R Soc Lond B*, 314: 1–340, 1986. [23](#)
- [201] J. A. White, J. T. Rubinstein, and A. R. Kay. Channel noise in neurons. *Trends Neurosci.*, 23(3):131–137, 2000. [99](#), [106](#), [107](#)
- [202] C. J. Wilson and Y. Kawaguchi. The origins of two-state spontaneous membrane potential fluctuations of neostriatal spiny neurons. *J Neurosci*, 16(7):2397–2410, 1996. [69](#)
- [203] C. J. Wong. Afferent and efferent connections of the diencephalic prepacemaker nucleus in the weakly electric fish, *Eigenmannia virescens*: interactions between the electromotor system and the neuroendocrine axis. *J Comp Neurol*, 383(1):18–41, 1997. [34](#), [49](#)
- [204] I. G. Woods, D. Schoppik, V. J. Shi, S. Zimmerman, H. A. Coleman, J. Greenwood, E. R. Soucy, and A. F. Schier. Neuropeptidergic signaling partitions arousal behaviors in zebrafish. *J Neurosci*, 34(9):3142–3160, 2014. [49](#), [68](#)
- [205] T. Yamamoto and L. Maler. Organization of galanin-like immunoreactive neuronal systems in weakly electric fish (*Apteronotus leptorhynchus*). *J Chem Neuroanat* 5:19–38. *J Chem Neuroanat*, 5(1):19–38, 1992. [68](#)
- [206] Y. Yarom and J. Hounsgaard. Voltage fluctuations in neurons: signal or noise? *Physiol Rev*, 91(3):917–929, 2011. [139](#)
- [207] M. M. Yartsev. The emperor’s new wardrobe: Rebalancing diversity of animal models in neuroscience research. *Science*, 358(6362):466–469, 2017. [26](#)
- [208] E. Yemini, T. Jucikas, L. J. Grundy, A. E. Brown, and W. R. Schafer. A database of *C. elegans* behavioral phenotypes. *Nat Methods*, 10(9):877–879, 2013. [23](#)

- [209] J. Zhang, C. M. Carver, F. S. Choveau, and M. S. Shapiro. Clustering and functional coupling of diverse ion channels and signaling proteins revealed by super-resolution STORM microscopy in neurons. *Neuron*, 92(2):461–478, 2016. [107](#)
- [210] G. K. H. Zupanc and L. Maler. Evoked chirping in the weakly electric fish *Apteronotus leptorhynchus*: a quantitative biophysical analysis. *Can. J. Zool.*, 71(11):2301–2310, 1993. [130](#)
- [211] G. K. H. Zupanc, R. F. Sîrbulescu, A. Nichols, and I. Ilies. Electric interactions through chirping behavior in the weakly electric fish, *Apteronotus leptorhynchus*. *J. Comp. Physiol. A Neuroethol. Sens. Neural. Behav. Physiol.*, 192(2):159–173, 2006. [28](#)